



Process Management Interface for Exascale (PMIx) Standard

Version 5.0 (Draft)

Created on September 19, 2022

This document describes the Process Management Interface for Exascale (PMIx) Standard, version 5.0 (Draft).

Comments: Please provide comments on the PMIx Standard by filing issues on the document repository <https://github.com/pmix/pmix-standard/issues> or by sending them to the PMIx Community mailing list at <https://groups.google.com/forum/#!forum/pmix>. Comments should include the version of the PMIx standard you are commenting about, and the page, section, and line numbers that you are referencing. Please note that messages sent to the mailing list from an unsubscribed e-mail address will be ignored.

Copyright © 2018-2020 PMIx Administrative Steering Committee (ASC).

Permission to copy without fee all or part of this material is granted, provided the PMIx ASC copyright notice and the title of this document appear, and notice is given that copying is by permission of PMIx ASC.

This page intentionally left blank

Unofficial Draft

Contents

1. Introduction	1
1.1. Background	1
1.2. PMIx Architecture Overview	1
1.3. Portability of Functionality	3
1.3.1. Attributes in PMIx	3
1.3.2. PMIx Roles	5
2. PMIx Terms and Conventions	7
2.1. Notational Conventions	9
2.2. Semantics	11
2.3. Naming Conventions	11
2.4. Procedure Conventions	11
3. Data Structures and Types	13
3.1. Constants	14
3.1.1. PMIx Return Status Constants	15
3.1.1.1. User-Defined Error and Event Constants	17
3.2. Data Types	17
3.2.1. Key Structure	17
3.2.1.1. Key support macros	18
3.2.2. Namespace Structure	19
3.2.2.1. Namespace support macros	19
3.2.3. Rank Structure	20
3.2.3.1. Rank support macros	21
3.2.4. Process Structure	21
3.2.4.1. Process structure support macros	22
3.2.5. Process State Structure	26
3.2.6. Process Information Structure	27
3.2.6.1. Process information structure support macros	27

3.2.7.	Job State Structure	29
3.2.8.	Value Structure	29
3.2.8.1.	Value structure support	30
3.2.9.	Info Structure	34
3.2.9.1.	Info structure support macros	34
3.2.9.2.	Info structure list macros	37
3.2.10.	Info Type Directives	40
3.2.10.1.	Info Directive support macros	41
3.2.11.	Environmental Variable Structure	43
3.2.11.1.	Environmental variable support macros	43
3.2.12.	Byte Object Type	45
3.2.12.1.	Byte object support macros	45
3.2.13.	Data Array Structure	47
3.2.13.1.	Data array support macros	47
3.2.14.	Argument Array Macros	49
3.2.15.	Set Environment Variable	53
3.3.	Generalized Data Types Used for Packing/Unpacking	53
3.4.	General Callback Functions	56
3.4.1.	Release Callback Function	56
3.4.2.	Lookup Callback Function	56
3.4.3.	Op Callback Function	57
3.4.4.	Value Callback Function	57
3.4.5.	Info Callback Function	58
3.4.6.	Handler registration callback function	58
3.5.	PMIx Datatype Value String Representations	59
4.	Client Initialization and Finalization	63
4.1.	PMIx_Initialized	63
4.2.	PMIx_Get_version	64
4.3.	PMIx_Init	64
4.3.1.	Initialization events	67
4.3.2.	Initialization attributes	67
4.3.2.1.	Connection attributes	67
4.3.2.2.	Programming model attributes	68

4.4.	PMIx_Finalize	69
4.4.1.	Finalize attributes	69
4.5.	PMIx_Progress	69
5.	Data Access and Sharing	71
5.1.	Non-reserved keys	73
5.2.	Posting Key/Value Pairs	74
5.2.1.	PMIx_Put	74
5.2.1.1.	Scope of Put Data	75
5.2.2.	PMIx_Store_internal	76
5.2.3.	PMIx_Commit	76
5.3.	Retrieval rules for non-reserved keys	77
5.4.	PMIx_Get	78
5.4.1.	PMIx_Get_nb	81
5.4.2.	Retrieval attributes	84
6.	Reserved Keys	86
6.1.	Data realms	86
6.1.1.	Session realm attributes	87
6.1.2.	Job realm attributes	89
6.1.3.	Application realm attributes	91
6.1.4.	Process realm attributes	92
6.1.5.	Node realm keys	94
6.2.	Retrieval rules for reserved keys	95
6.2.1.	Accessing information: examples	96
6.2.1.1.	Session-level information	96
6.2.1.2.	Job-level information	97
6.2.1.3.	Application-level information	98
6.2.1.4.	Process-level information	99
6.2.1.5.	Node-level information	99
7.	Query Operations	100
7.1.	PMIx_Query_info	100
7.1.1.	Query Structure	101
7.1.2.	PMIx_Query_info	101

7.1.3.	PMIx_Query_info_nb	106
7.1.4.	Query keys	110
7.1.5.	Query attributes	112
7.1.5.1.	Query structure support macros	113
7.2.	PMIx_Resolve_peers	115
7.2.1.	PMIx_Resolve_nodes	116
7.3.	Using Get vs Query	117
7.4.	Accessing attribute support information	118
8.	Synchronization	120
8.1.	PMIx_Fence	120
8.2.	PMIx_Fence_nb	122
8.2.1.	Fence-related attributes	124
9.	Publish/Lookup Operations	126
9.1.	PMIx_Publish2	126
9.2.	PMIx_Publish2_nb	128
9.3.	PMIx_Publish	130
9.4.	PMIx_Publish_nb	131
9.5.	Publish-specific constants	133
9.6.	Publish-specific attributes	133
9.7.	Publish-Lookup Datatypes	133
9.7.1.	Range of Published Data	134
9.7.2.	Data Persistence Structure	134
9.7.3.	Lookup Related Data Structures	135
9.8.	PMIx_Lookup	135
9.9.	PMIx_Lookup_nb	137
9.9.0.1.	Lookup data structure support macros	139
9.10.	Retrieval rules for published data	142
9.11.	PMIx_Unpublish	142
9.12.	PMIx_Unpublish_nb	143
10.	Event Notification	146
10.1.	Notification and Management	146
10.1.1.	Events versus status constants	148

10.1.2.	PMIx_Register_event_handler	148
10.1.3.	Event registration constants	151
10.1.4.	System events	152
10.1.5.	Event handler registration and notification attributes	152
10.1.5.1.	Fault tolerance event attributes	153
10.1.5.2.	Hybrid programming event attributes	153
10.1.6.	Notification Function	154
10.1.7.	PMIx_Deregister_event_handler	155
10.1.8.	PMIx_Notify_event	156
10.1.9.	Notification Handler Completion Callback Function	160
10.1.9.1.	Completion Callback Function Status Codes	160
11.	Data Packing and Unpacking	161
11.1.	Data Buffer Type	161
11.2.	Support Macros	162
11.3.	General Routines	164
11.3.1.	PMIx_Data_pack	164
11.3.2.	PMIx_Data_unpack	165
11.3.3.	PMIx_Data_copy	167
11.3.4.	PMIx_Data_print	168
11.3.5.	PMIx_Data_copy_payload	168
11.3.6.	PMIx_Data_load	169
11.3.7.	PMIx_Data_unload	170
11.3.8.	PMIx_Data_compress	171
11.3.9.	PMIx_Data_decompress	172
11.3.10.	PMIx_Data_embed	173
12.	Process Management	174
12.1.	Abort	174
12.1.1.	PMIx_Abort	174
12.2.	Process Creation	175
12.2.1.	PMIx_Spawn	176
12.2.2.	PMIx_Spawn_nb	181
12.2.3.	Spawn-specific constants	187

12.2.4.	Spawn attributes	188
12.2.5.	Application Structure	192
12.2.5.1.	App structure support macros	192
12.2.5.2.	Spawn Callback Function	194
12.3.	Connecting and Disconnecting Processes	195
12.3.1.	PMIx_Connect	195
12.3.2.	PMIx_Connect_nb	197
12.3.3.	PMIx_Disconnect	199
12.3.4.	PMIx_Disconnect_nb	201
12.4.	Process Locality	203
12.4.1.	PMIx_Load_topology	203
12.4.2.	PMIx_Get_relative_locality	204
12.4.2.1.	Topology description	204
12.4.2.2.	Topology support macros	205
12.4.2.3.	Relative locality of two processes	206
12.4.2.4.	Locality keys	206
12.4.3.	PMIx_Parse_cpuset_string	206
12.4.4.	PMIx_Get_cpuset	207
12.4.4.1.	Binding envelope	208
12.4.5.	PMIx_Compute_distances	208
12.4.6.	PMIx_Compute_distances_nb	209
12.4.7.	Device Distance Callback Function	210
12.4.8.	Device type	211
12.4.9.	Device Distance Structure	211
12.4.10.	Device distance support macros	212
12.4.11.	Device distance attributes	214
13.	Job Management and Reporting	215
13.1.	Allocation Requests	215
13.1.1.	PMIx_Allocation_request	215
13.1.2.	PMIx_Allocation_request_nb	218
13.1.3.	Job Allocation attributes	221
13.1.4.	Job Allocation Directives	223

13.2.	Job Control	223
13.2.1.	PMIx_Job_control	224
13.2.2.	PMIx_Job_control_nb	226
13.2.3.	Job control constants	229
13.2.4.	Job control events	229
13.2.5.	Job control attributes	230
13.3.	Process and Job Monitoring	231
13.3.1.	PMIx_Process_monitor	231
13.3.2.	PMIx_Process_monitor_nb	233
13.3.3.	PMIx_Heartbeat	235
13.3.4.	Monitoring events	236
13.3.5.	Monitoring attributes	236
13.4.	Logging	237
13.4.1.	PMIx_Log	237
13.4.2.	PMIx_Log_nb	240
13.4.3.	Log attributes	243
14.	Process Sets and Groups	245
14.1.	Process Sets	245
14.1.1.	Process Set Constants	246
14.1.2.	Process Set Attributes	247
14.2.	Process Groups	247
14.2.1.	Relation to the host environment	247
14.2.2.	Construction procedure	248
14.2.3.	Destruct procedure	249
14.2.4.	Process Group Events	249
14.2.5.	Process Group Attributes	250
14.2.6.	PMIx_Group_construct	252
14.2.7.	PMIx_Group_construct_nb	255
14.2.8.	PMIx_Group_destruct	258
14.2.9.	PMIx_Group_destruct_nb	259
14.2.10.	PMIx_Group_invite	261
14.2.11.	PMIx_Group_invite_nb	264
14.2.12.	PMIx_Group_join	266

14.2.13. PMIx_Group_join_nb	268
14.2.13.1. Group accept/decline directives	270
14.2.14. PMIx_Group_leave	270
14.2.15. PMIx_Group_leave_nb	271
15. Fabric Support Definitions	273
15.1. Fabric Support Events	276
15.2. Fabric Support Datatypes	276
15.2.1. Fabric Endpoint Structure	276
15.2.2. Fabric endpoint support macros	277
15.2.3. Fabric Coordinate Structure	278
15.2.4. Fabric coordinate support macros	279
15.2.5. Fabric Geometry Structure	280
15.2.6. Fabric geometry support macros	280
15.2.7. Fabric Coordinate Views	282
15.2.8. Fabric Link State	282
15.2.9. Fabric Operation Constants	283
15.2.10. Fabric registration structure	283
15.2.10.1. Static initializer for the fabric structure	286
15.2.10.2. Initialize the fabric structure	286
15.3. Fabric Support Attributes	287
15.4. Fabric Support Functions	290
15.4.1. PMIx_Fabric_register	290
15.4.2. PMIx_Fabric_register_nb	292
15.4.3. PMIx_Fabric_update	293
15.4.4. PMIx_Fabric_update_nb	294
15.4.5. PMIx_Fabric_deregister	294
15.4.6. PMIx_Fabric_deregister_nb	295
16. Security	296
16.1. Obtaining Credentials	296
16.1.1. PMIx_Get_credential	297
16.1.2. PMIx_Get_credential_nb	298
16.1.3. Credential Attributes	299

16.2. Validating Credentials	300
16.2.1. PMIx_Validate_credential	300
16.2.2. PMIx_Validate_credential_nb	301
17. Server-Specific Interfaces	304
17.1. Server Initialization and Finalization	304
17.1.1. PMIx_server_init	304
17.1.2. PMIx_server_finalize	308
17.1.3. Server Initialization Attributes	309
17.2. Server Support Functions	310
17.2.1. PMIx_generate_regex	310
17.2.2. PMIx_generate_ppn	311
17.2.3. PMIx_server_register_namespace	312
17.2.3.1. Namespace registration attributes	322
17.2.3.2. Assembling the registration information	323
17.2.4. PMIx_server_deregister_namespace	331
17.2.5. PMIx_server_register_resources	333
17.2.6. PMIx_server_deregister_resources	334
17.2.7. PMIx_server_register_client	335
17.2.8. PMIx_server_deregister_client	336
17.2.9. PMIx_server_setup_fork	337
17.2.10. PMIx_server_dmodex_request	337
17.2.10.1. Server Direct Modex Response Callback Function	339
17.2.11. PMIx_server_setup_application	339
17.2.11.1. Server Setup Application Callback Function	343
17.2.11.2. Server Setup Application Attributes	343
17.2.12. PMIx_Register_attributes	344
17.2.12.1. Attribute registration constants	345
17.2.12.2. Attribute registration structure	345
17.2.12.3. Attribute registration structure descriptive attributes	346
17.2.12.4. Attribute registration structure support macros	346
17.2.13. PMIx_server_setup_local_support	348
17.2.14. PMIx_server_IOF_deliver	350
17.2.15. PMIx_server_collect_inventory	351

17.2.16.	PMIx_server_deliver_inventory	352
17.2.17.	PMIx_server_generate_locality_string	353
17.2.18.	PMIx_server_generate_cpuset_string	354
17.2.18.1.	Cpuset Structure	355
17.2.18.2.	Cpuset support macros	355
17.2.19.	PMIx_server_define_process_set	356
17.2.20.	PMIx_server_delete_process_set	357
17.3.	Server Function Pointers	358
17.3.1.	pmix_server_module_t Module	358
17.3.2.	pmix_server_client_connected_fn_t	360
17.3.3.	pmix_server_client_connected2_fn_t	361
17.3.4.	pmix_server_client_finalized_fn_t	362
17.3.5.	pmix_server_abort_fn_t	363
17.3.6.	pmix_server_fencefn_fn_t	365
17.3.6.1.	Modex Callback Function	368
17.3.7.	pmix_server_dmodex_req_fn_t	368
17.3.7.1.	Dmodex attributes	370
17.3.8.	pmix_server_publish_fn_t	370
17.3.9.	pmix_server_lookup_fn_t	372
17.3.10.	pmix_server_unpublish_fn_t	375
17.3.11.	pmix_server_spawn_fn_t	377
17.3.11.1.	Server spawn attributes	382
17.3.12.	pmix_server_connect_fn_t	382
17.3.13.	pmix_server_disconnect_fn_t	384
17.3.14.	pmix_server_register_events_fn_t	387
17.3.15.	pmix_server_deregister_events_fn_t	388
17.3.16.	pmix_server_notify_event_fn_t	390
17.3.17.	pmix_server_listener_fn_t	391
17.3.17.1.	PMIx Client Connection Callback Function	392
17.3.18.	pmix_server_query_fn_t	393
17.3.19.	pmix_server_tool_connection_fn_t	395
17.3.19.1.	Tool connection attributes	398
17.3.19.2.	PMIx Tool Connection Callback Function	398

17.3.20.	<code>pmix_server_log_fn_t</code>	398
17.3.21.	<code>pmix_server_alloc_fn_t</code>	400
17.3.22.	<code>pmix_server_job_control_fn_t</code>	403
17.3.23.	<code>pmix_server_monitor_fn_t</code>	406
17.3.24.	<code>pmix_server_get_cred_fn_t</code>	409
17.3.24.1.	Credential callback function	410
17.3.25.	<code>pmix_server_validate_cred_fn_t</code>	411
17.3.26.	Credential validation callback function	413
17.3.27.	<code>pmix_server_iof_fn_t</code>	414
17.3.27.1.	IOF delivery function	417
17.3.28.	<code>pmix_server_stdin_fn_t</code>	418
17.3.29.	<code>pmix_server_grp_fn_t</code>	419
17.3.29.1.	Group Operation Constants	422
17.3.30.	<code>pmix_server_fabric_fn_t</code>	422
18.	Tools and Debuggers	425
18.1.	Connection Mechanisms	425
18.1.1.	Rendezvousing with a local server	428
18.1.2.	Connecting to a remote server	429
18.1.3.	Attaching to running jobs	430
18.1.4.	Tool initialization attributes	430
18.1.5.	Tool initialization environmental variables	431
18.1.6.	Tool connection attributes	431
18.2.	Launching Applications with Tools	432
18.2.1.	Direct launch	432
18.2.2.	Indirect launch	438
18.2.2.1.	Initiator-based command line parsing	438
18.2.2.2.	Intermediate Launcher (IL)-based command line parsing	441
18.2.3.	Tool spawn-related attributes	442
18.2.4.	Tool rendezvous-related events	443
18.3.	IO Forwarding	444
18.3.1.	Forwarding stdout/stderr	444
18.3.2.	Forwarding stdin	447
18.3.3.	IO Forwarding Channels	448

18.3.4.	IO Forwarding constants	448
18.3.5.	IO Forwarding attributes	449
18.4.	Debugger Support	450
18.4.1.	Co-Location of Debugger Daemons	453
18.4.2.	Co-Spawn of Debugger Daemons	454
18.4.3.	Debugger Agents	455
18.4.4.	Tracking the job lifecycle	456
18.4.4.1.	Job lifecycle events	457
18.4.4.2.	Job lifecycle attributes	458
18.4.5.	Debugger-related constants	458
18.4.6.	Debugger attributes	459
18.5.	Tool-Specific APIs	460
18.5.1.	PMIx_tool_init	461
18.5.2.	PMIx_tool_finalize	463
18.5.3.	PMIx_tool_disconnect	464
18.5.4.	PMIx_tool_attach_to_server	465
18.5.5.	PMIx_tool_get_servers	466
18.5.6.	PMIx_tool_set_server	467
18.5.7.	PMIx_IOF_pull	468
18.5.8.	PMIx_IOF_deregister	470
18.5.9.	PMIx_IOF_push	471
19.	Storage Support Definitions	474
19.1.	Storage support constants	474
19.2.	Storage support attributes	476
A.	Python Bindings	478
A.1.	Design Considerations	478
A.1.1.	Error Codes vs Python Exceptions	478
A.1.2.	Representation of Structured Data	478
A.2.	Datatype Definitions	479
A.2.1.	Example	485
A.3.	Callback Function Definitions	486
A.3.1.	IOF Delivery Function	486

A.3.2.	Event Handler	486
A.3.3.	Server Module Functions	487
A.3.3.1.	Client Connected	487
A.3.3.2.	Client Finalized	488
A.3.3.3.	Client Aborted	488
A.3.3.4.	Fence	489
A.3.3.5.	Direct Modex	490
A.3.3.6.	Publish	490
A.3.3.7.	Lookup	491
A.3.3.8.	Unpublish	491
A.3.3.9.	Spawn	492
A.3.3.10.	Connect	492
A.3.3.11.	Disconnect	493
A.3.3.12.	Register Events	493
A.3.3.13.	Deregister Events	494
A.3.3.14.	Notify Event	494
A.3.3.15.	Query	494
A.3.3.16.	Tool Connected	495
A.3.3.17.	Log	495
A.3.3.18.	Allocate Resources	496
A.3.3.19.	Job Control	496
A.3.3.20.	Monitor	497
A.3.3.21.	Get Credential	497
A.3.3.22.	Validate Credential	498
A.3.3.23.	IO Forward	498
A.3.3.24.	IO Push	499
A.3.3.25.	Group Operations	499
A.3.3.26.	Fabric Operations	500
A.4.	PMIxCClient	501
A.4.1.	Client.init	501
A.4.2.	Client.initialized	501
A.4.3.	Client.get_version	502
A.4.4.	Client.finalize	502

A.4.5.	Client.abort	502
A.4.6.	Client.store_internal	503
A.4.7.	Client.put	503
A.4.8.	Client.commit	504
A.4.9.	Client.fence	504
A.4.10.	Client.get	505
A.4.11.	Client.publish	505
A.4.12.	Client.lookup	506
A.4.13.	Client.unpublish	506
A.4.14.	Client.spawn	507
A.4.15.	Client.connect	507
A.4.16.	Client.disconnect	508
A.4.17.	Client.resolve_peers	508
A.4.18.	Client.resolve_nodes	509
A.4.19.	Client.query	509
A.4.20.	Client.log	510
A.4.21.	Client.allocation_request	510
A.4.22.	Client.job_ctrl	511
A.4.23.	Client.monitor	511
A.4.24.	Client.get_credential	512
A.4.25.	Client.validate_credential	512
A.4.26.	Client.group_construct	513
A.4.27.	Client.group_inyite	513
A.4.28.	Client.group_join	514
A.4.29.	Client.group_leave	515
A.4.30.	Client.group_destruct	515
A.4.31.	Client.register_event_handler	515
A.4.32.	Client.deregister_event_handler	516
A.4.33.	Client.notify_event	516
A.4.34.	Client.fabric_register	517
A.4.35.	Client.fabric_update	517
A.4.36.	Client.fabric_deregister	518
A.4.37.	Client.load_topology	518

A.4.38.	Client.get_relative_locality	519
A.4.39.	Client.get_cpuset	519
A.4.40.	Client.parse_cpuset_string	519
A.4.41.	Client.compute_distances	520
A.4.42.	Client.error_string	520
A.4.43.	Client.proc_state_string	521
A.4.44.	Client.scope_string	521
A.4.45.	Client.persistence_string	522
A.4.46.	Client.data_range_string	522
A.4.47.	Client.info_directives_string	522
A.4.48.	Client.data_type_string	523
A.4.49.	Client.alloc_directive_string	523
A.4.50.	Client.iof_channel_string	524
A.4.51.	Client.job_state_string	524
A.4.52.	Client.get_attribute_string	524
A.4.53.	Client.get_attribute_name	525
A.4.54.	Client.link_state_string	525
A.4.55.	Client.device_type_string	526
A.4.56.	Client.progress	526
A.5.	PMIxServer	526
A.5.1.	Server.init	526
A.5.2.	Server.finalize	527
A.5.3.	Server.generate_regex	527
A.5.4.	Server.generate_ppn	528
A.5.5.	Server.generate_locality_string	528
A.5.6.	Server.generate_cpuset_string	529
A.5.7.	Server.register_nspace	529
A.5.8.	Server.deregister_nspace	530
A.5.9.	Server.register_resources	530
A.5.10.	Server.deregister_resources	531
A.5.11.	Server.register_client	531
A.5.12.	Server.deregister_client	532
A.5.13.	Server.setup_fork	532

A.5.14.	Server.dmodex_request	532
A.5.15.	Server.setup_application	533
A.5.16.	Server.register_attributes	533
A.5.17.	Server.setup_local_support	534
A.5.18.	Server.iof_deliver	534
A.5.19.	Server.collect_inventory	535
A.5.20.	Server.deliver_inventory	535
A.5.21.	Server.define_process_set	536
A.5.22.	Server.delete_process_set	536
A.5.23.	Server.register_resources	537
A.5.24.	Server.deregister_resources	537
A.6.	PMIxTool	538
A.6.1.	Tool.init	538
A.6.2.	Tool.finalize	538
A.6.3.	Tool.disconnect	538
A.6.4.	Tool.attach_to_server	539
A.6.5.	Tool.get_servers	539
A.6.6.	Tool.set_server	540
A.6.7.	Tool.iof_pull	540
A.6.8.	Tool.iof_deregister	541
A.6.9.	Tool.iof_push	541
A.7.	Example Usage	542
A.7.1.	Python Client	542
A.7.2.	Python Server	544
B.	Use-Cases	548
B.1.	Business Card Exchange for Process-to-Process Wire-up	548
B.1.1.	Use Case Summary	548
B.1.2.	Use Case Details	549
B.2.	Debugging	552
B.2.1.	Terminology	552
B.2.1.1.	Tools vs Debuggers	552
B.2.1.2.	Parallel Launching Methods	553
B.2.1.3.	Process Synchronization	553

B.2.1.4.	Process Acquisition	553
B.2.2.	Use Case Details	553
B.2.2.1.	Direct-Launch Debugger Tool	553
B.2.2.2.	Indirect-Launch Debugger Tool	558
B.2.2.3.	Attaching to a Running Job	564
B.2.2.4.	Tool Interaction with RM	567
B.2.2.5.	Environmental Parameter Directives for Applications and Launchers	569
B.3.	Hybrid Applications	570
B.3.1.	Use Case Summary	570
B.3.2.	Use Case Details	570
B.3.2.1.	Identifying Active Parallel Runtime Systems	570
B.3.2.2.	Coordinating at Runtime	572
B.3.2.3.	Coordinating at Runtime with Multiple Event Handlers	574
B.4.	MPI Sessions	577
B.4.1.	Use Case Summary	577
B.4.2.	Use Case Details	578
B.5.	Cross-Version Compatibility	580
B.5.1.	Use Case Summary	580
B.5.2.	Use Case Details	580
C.	Revision History	584
C.1.	Version 1.0: June 12, 2015	584
C.2.	Version 2.0: Sept. 2018	585
C.2.1.	Removed/Modified Application Programming Interfaces (APIs)	585
C.2.2.	Deprecated constants	585
C.2.3.	Deprecated attributes	586
C.3.	Version 2.1: Dec. 2018	586
C.4.	Version 2.2: Jan 2019	587
C.5.	Version 3.0: Dec. 2018	587
C.5.1.	Removed constants	588
C.5.2.	Deprecated attributes	588
C.5.3.	Removed attributes	588
C.6.	Version 3.1: Jan. 2019	589

C.7.	Version 3.2: Oct. 2020	589
C.7.1.	Deprecated constants	590
C.7.2.	Deprecated attributes	591
C.8.	Version 4.0: Dec. 2020	592
C.8.1.	Added Constants	594
C.8.2.	Added Attributes	597
C.8.3.	Added Environmental Variables	610
C.8.4.	Added Macros	610
C.8.5.	Deprecated APIs	611
C.8.6.	Deprecated constants	611
C.8.7.	Removed constants	611
C.8.8.	Deprecated attributes	612
C.8.9.	Removed attributes	613
C.9.	Version 4.1: TBD	614
C.9.1.	Removed constants	614
C.9.2.	Added Functions (Provisional)	614
C.9.3.	Added Data Structures (Provisional)	615
C.9.4.	Added Macros (Provisional)	615
C.9.5.	Added Constants (Provisional)	615
C.9.6.	Added Attributes (Provisional)	616
C.10.	Version 4.2: TBD	617
C.10.1.	Deprecated constants	617
C.10.2.	Deprecated attributes	617
C.10.3.	Deprecated macros	618
C.10.4.	Added Functions (Provisional)	618
C.10.5.	Added Macros (Provisional)	618
C.10.6.	Added Constants (Provisional)	619
C.10.7.	Added Attributes (Provisional)	619
D.	Acknowledgements	622
D.1.	Version 4.0	622
D.2.	Version 3.0	623
D.3.	Version 2.0	624
D.4.	Version 1.0	625

Bibliography	627
Index	628
Index of APIs	630
Index of Support Macros	638
Index of Data Structures	643
Index of Constants	645
Index of Environmental Variables	655
Index of Attributes	656

Unofficial Draft

CHAPTER 1

Introduction

1 Process Management Interface - Exascale (PMIx) is an application programming interface standard
2 that provides libraries and programming models with portable and well-defined access to commonly
3 needed services in distributed and parallel computing systems. A typical example of such a service
4 is the portable and scalable exchange of network addresses to establish communication channels
5 between the processes of a parallel application or service. As such, PMIx gives distributed system
6 software providers a better understanding of how programming models and libraries can interface
7 with and use system-level services. As a standard, PMIx provides APIs that allow for portable
8 access to these varied system software services and the functionalities they offer. Although these
9 services can be defined and implemented directly by the system software components providing
10 them, the community represented by the ASC feels that the development of a shared standard better
11 serves the community. As a result, PMIx enables programming languages and libraries to focus on
12 their core competencies without having to provide their own system-level services.

1.1 Background

14 The Process Management Interface (PMI) has been used for quite some time as a means of
15 exchanging wireup information needed for inter-process communication. Two versions (PMI-1 and
16 PMI-2 [2]) have been released as part of the MPICH effort, with PMI-2 demonstrating better
17 scaling properties than its PMI-1 predecessor.

18 PMI-1 and PMI-2 can be implemented using PMIx though PMIx is not a strict superset of either.
19 Since its introduction, PMIx has expanded on earlier PMI efforts by providing an extended version
20 of the PMI APIs which provide necessary functionality for launching and managing parallel
21 applications and tools at scale.

22 The increase in adoption has motivated the creation of this document to formally specify the
23 intended behavior of the PMIx APIs.

24 More information about the PMIx standard and affiliated projects can be found at the PMIx web
25 site: <https://pmix.org>

1.2 PMIx Architecture Overview

27 The presentation of the PMIx APIs within this document makes some basic assumptions about how
28 these APIs are used and implemented. These assumptions are generally made only to simplify the
29 presentation and explain PMIx with the expectation that most readers have similar concepts on how

1 computing systems are organized today. However, ultimately this document should only be
2 assumed to define a set of APIs.

3 A concept that is fundamental to PMIx is that a PMIx implementation might operate primarily as a
4 *messenger*, and not a *doer* — i.e., a PMIx implementation might rely heavily or fully on other
5 software components to provide functionality [1]. Since a PMIx implementation might only deliver
6 requests and responses to other software components, the API calls include ways to provide
7 arbitrary information to the backend components that actually implement the functionality. Also,
8 because PMIx implementations generally rely heavily on other system software, a PMIx
9 implementation might not be able to guarantee that a feature is available on all platforms the
10 implementation supports. These aspects are discussed in detail in the remainder of this chapter.

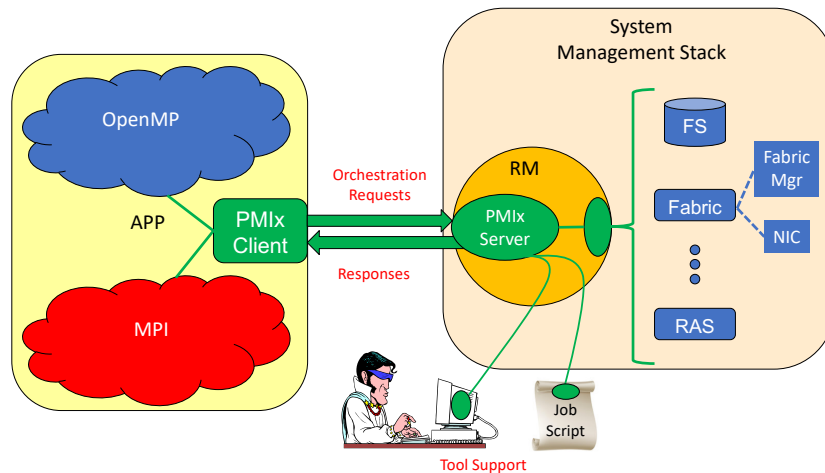


Figure 1.1.: PMIx-SMS Interactions

11 Fig. 1.1 shows a typical PMIx implementation in which the application is built against a PMIx
12 client library that contains the client-side APIs, attribute definitions, and communication support
13 for interacting with the local PMIx server. PMIx clients are processes which are started through the
14 PMIx infrastructure, either by the PMIx implementation directly or through a System Management
15 Software stack (SMS) component, and have registered as clients. A PMIx client is created in such a
16 way that the PMIx client library will have sufficient information available to authenticate with
17 the PMIx server. The PMIx server will have sufficient knowledge about the process which it
18 created, either directly or through other SMS, to authenticate the process and provide information
19 the process requests such as its identity and the identity of its peers.

20 As clients invoke PMIx APIs, it is possible that some client requests can be handled at the client
21 level. Other requests might require communication with the local PMIx server, which subsequently
22 might request services from the host SMS (represented here by a Resource Manager (RM)
23 daemon). The interaction between the PMIx server and SMS are achieved using callback functions
24 registered during server initialization. The host SMS can indicate its lack of support for any

1 operation by simply providing a *NULL* for the associated callback function, or can create a function
2 entry that returns *not supported* when called.

3 Recognizing the burden this places on SMS vendors, the PMIx community has included interfaces
4 by which the host SMS (containing the local PMIx service instance) can request support from local
5 SMS elements via the PMIx API. Once the SMS has transferred the request to an appropriate
6 location, a PMIx server interface can be used to pass the request between SMS subsystems. For
7 example, a request for network traffic statistics can utilize the PMIx networking abstractions to
8 retrieve the information from the Fabric Manager. This reduces the portability and interoperability
9 issues between the individual subsystems by transferring the burden of defining the interoperable
10 interfaces from the SMS subsystems to the PMIx community, which continues to work with those
11 providers to develop the necessary support.

12 Fig. 1.1 shows how tools can interact with the PMIx architecture. Tools, whether standalone or
13 embedded in job scripts, are an exception to the normal client registration process. A process can
14 register as a tool, provided the PMIx client library has adequate rendezvous information to connect
15 to the appropriate PMIx server (either hosted on the local machine or on a remote machine). This
16 allows processes which were not created by the PMIx infrastructure to request access to PMIx
17 functionality.

18 1.3 Portability of Functionality

19 It is difficult to define a portable API that will provide access to the many and varied features
20 underlying the operations for which PMIx provides access. For example, the options and features
21 provided to request the creation of new processes varied dramatically between different systems
22 existing at the time PMIx was introduced. Many RMs provide rich interfaces to specify the
23 resources assigned to processes. As a result, PMIx is faced with the challenge of attempting to meet
24 the seemingly conflicting goals of creating an API which allows access to these diverse features
25 while being portable across a wide range of existing software environments. In addition, the
26 functionalities required by different clients vary greatly. Producing a PMIx implementation which
27 can provide the needs of all possible clients on all of its target systems could be so burdensome as
28 to discourage PMIx implementations.

29 To help address this issue, the PMIx APIs are designed to allow resource managers and other
30 system management stack components to decide on support of a particular function and allow client
31 applications to query and adjust to the level of support available. PMIx clients should be written to
32 account for the possibility that a PMIx API might return an error code indicating that the call is not
33 supported. The PMIx community continues to look at ways to assist SMS implementers in their
34 decisions on what functionality to support by highlighting functions and attributes that are critical
35 to basic application execution (e.g., [PMIx_Get](#)) for certain classes of applications.

36 1.3.1 Attributes in PMIx

37 An area where differences between support on different systems can be challenging is regarding the
38 attributes that provide information to the client process and/or control the behavior of a PMIx API.

1 Most PMIx API calls can accept additional information or attributes specified in the form of
2 key/value pairs. These attributes provide information to the PMIx implementation that influence the
3 behavior of the API call. In addition to API calls being optional, support for the individual
4 attributes of an API call can vary between systems or implementations.

5 An application can adapt to the attribute support on a particular system in one of two ways. PMIx
6 provides an API to enable an application to query the attributes supported by a particular API (See
7 7.4). Through this API, the PMIx implementation can provide detailed information about the
8 attributes supported on a system for each API call queried. Alternatively, the application can mark
9 attributes as required using a flag within the `pmix_info_t` (See 3.2.9). If the required attribute is
10 not available on the system or the desired value for the attribute is not available, the call will return
11 the error code for *not supported*.

12 For example, the `PMIX_TIMEOUT` attribute can be used to specify the time (in seconds) before the
13 requested operation should time out. The intent of this attribute is to allow the client to avoid
14 “hanging” in a request that takes longer than the client wishes to wait, or may never return (e.g., a
15 `PMIx_Fence` that a blocked participant never enters).

16 The application can query the attribute support for `PMIx_Fence` and search whether
17 `PMIX_TIMEOUT` is listed as a supported attribute. The application can also set the required flag in
18 the `pmix_info_t` for that attribute when making the `PMIx_Fence` call. This will return an
19 error if this attribute is not supported. If the required flag is not set, the library and SMS host are
20 allowed to treat the attribute as optional, ignoring it if support is not available.

21 It is therefore critical that users and application implementers:

- 22 a) consider whether or not a given attribute is required, marking it accordingly; and
- 23 b) check the return status on all PMIx function calls to ensure support was present and that the
24 request was accepted. Note that for non-blocking APIs, a return of `PMIX_SUCCESS` only
25 indicates that the request had no obvious errors and is being processed – the eventual callback
26 will return the status of the requested operation itself.

27 PMIx clients (e.g., tools, parallel programming libraries) may find that they depend only on a small
28 subset of interfaces and attributes to work correctly. PMIx clients are strongly advised to define a
29 document itemizing the PMIx interfaces and associated attributes that are required for correct
30 operation, and are optional but recommended for full functionality. The PMIx standard cannot
31 define this list for all given PMIx clients, but such a list is valuable to RMs desiring to support these
32 clients.

33 A PMIx implementation may be able to support only a subset of the PMIx API and attributes on a
34 particular system due to either its own limitations or limitations of the SMS with which it
35 interfaces. A PMIx implementation may also provide additional attributes beyond those defined
36 herein in order to allow applications to access the full features of the underlying SMS. PMIx
37 implementations are strongly advised to document the PMIx interfaces and associated attributes
38 they support, with any annotations about behavior limitations. The PMIx standard cannot define
39 this support for implementations, but such documentation is valuable to PMIx clients desiring to
40 support a broad range of systems.

1 While a PMIx library implementer, or an SMS component server, may choose to support a
2 particular PMIx API, they are not required to support every attribute that might apply to it. This
3 would pose a significant barrier to entry for an implementer as there can be a broad range of
4 applicable attributes to a given API, at least some of which may rarely be used.

5 Note that an environment that does not include support for a particular attribute/API pair is not
6 “incomplete” or of lower quality than one that does include that support. Vendors must decide
7 where to invest their time based on the needs of their target markets, and it is perfectly reasonable
8 for them to perform cost/benefit decisions when considering what functions and attributes to
9 support.

10 Attributes in this document are organized according to their primary usage, either grouped with a
11 specific API or included in an appropriate functional chapter. Attributes in the PMIx Standard all
12 start with "**PMIX**" in their name, and many include a functional description as part of their name
13 (e.g., the use of "**PMIX_FABRIC_**" at the beginning of fabric-specific attributes). The PMIx
14 Standard also defines an attribute that can be used to indicate that an attribute variable has not yet
15 been set:

16 **PMIX_ATTR_UNDEF** "`pmix.undef`" (NULL)

17 A default attribute name signifying that the attribute field of a PMIx structure (e.g., a
18 `pmix_info_t`) has not yet been defined.

19 1.3.2 PMIx Roles

20 The role of a PMIx process in the PMIx universe is grouped into one of three categories based on
21 how it operates in the PMIx environment namely as a *client*, *server*, or *tool*. As a result, there are
22 three corresponding groupings of APIs each with their own initialization and finalization functions.
23 If a process initializes as either a *server* or a *tool* that process may also access all of the *client* APIs.

24 A process operating as a *client* is connected to the PMIx server instance within an RM when the
25 client calls the client PMIx initialization routine. The *client* is typically started directly or indirectly
26 (for example, by an intermediate script) by that RM. Additionally, a *client* may be started directly
27 by the user and then connect to an RM which is typically referred to as a *singleton* launch. A
28 process operating as a *server* is responsible for starting client processes and coordinating with other
29 server and tool processes in the same PMIx universe. Often processes operating as a *server* are part
30 of the Resource Manager (RM) infrastructure. A process operating as a *tool* is started
31 independently (e.g., via `fork/exec`) or by the RM and will connect to a PMIx *server* to interact with
32 the processes in the PMIx universe. An example of a *tool* process is a parallel debugger that will
33 connect to the server to assist with attaching to a set of client processes.

34 PMIx serves as a conduit between processes acting in these three different roles. As such, an API is
35 often described by how it interacts with processes operating in other roles in the PMIx universe.

Advice to PMIx library implementers

1 A PMIx implementation may support all or a subset of the API role groupings defined in the
2 standard. A common nomenclature is defined here to aid in identifying levels of conformance of an
3 implementation.

4 Note that it would not make sense for an implementation to exclude the *client* interfaces from their
5 implementation since they are also used by the *server* and *tool* roles. Therefore the *client* interfaces
6 represent the minimal set of required functionality for PMIx compliance.

7 A PMIx implementation that supports only the *client* APIs is said to be *client-role PMIx standard*
8 *compliant*. Similarly, a PMIx implementation that only supports the *client* and *tool* APIs is said to
9 be *client-role and tool-role PMIx standard compliant*. Finally, a PMIx implementation that only
10 supports the *client* and *server* APIs is said to be *client-role and server-role PMIx standard*
11 *compliant*.

12 A PMIx implementation that supports all three sets of the API role groupings is said to be
13 *client-role, server-role, and tool-role PMIx standard compliant*. These *client-role, server-role, and*
14 *tool-role PMIx standard compliant* implementations have the advantage of being able to support a
15 broad set of PMIx consumers in the different roles.

CHAPTER 2

PMIx Terms and Conventions

1 In this chapter we describe some common terms and conventions used throughout this document.
2 The PMIx Standard has adopted the widespread use of key-value *attributes* to add flexibility to the
3 functionality expressed in the APIs. Accordingly, the ASC has chosen to require that the definition
4 of each standard API include the passing of an array of attributes. These provide a means of
5 customizing the behavior of the API as future needs emerge without having to alter or create new
6 variants of it. In addition, attributes provide a mechanism by which researchers can easily explore
7 new approaches to a given operation without having to modify the API itself.

8 In an effort to maintain long-term backward compatibility, PMIx does not include large numbers of
9 APIs that each focus on a narrow scope of functionality, but instead relies on the definition of fewer
10 generic APIs that include arrays of key-value attributes for “tuning” the function’s behavior. Thus,
11 modifications to the PMIx standard primarily consist of the definition of new attributes along with a
12 description of the APIs to which they relate and the expected behavior when used with those APIs.

13 The following terminology is used throughout this document:

- 14 • *session* refers to a set of resources assigned by the WorkLoad Manager (WLM) that has been
15 reserved for one or more users. A session is identified by a *session ID* that is unique within the
16 scope of the governing WLMs. Historically, High Performance Computing (HPC) sessions have
17 consisted of a static allocation of resources - i.e., a block of resources assigned to a user in
18 response to a specific request and managed as a unified collection. However, this is changing in
19 response to the growing use of dynamic programming models that require on-the-fly allocation
20 and release of system resources. Accordingly, the term *session* in this document refers to a
21 potentially dynamic entity, perhaps comprised of resources accumulated as a result of multiple
22 allocation requests that are managed as a single unit by the WLM.
- 23 • *job* refers to a set of one or more *applications* executed as a single invocation by the user within a
24 session with a unique identifier, the *job ID*, assigned by the RM or launcher. For example, the
25 command line “*mpiexec -n 1 app1 : -n 2 app2*” generates a single Multiple Program Multiple
26 Data (MPMD) job containing two applications. A user may execute multiple *jobs* within a given
27 session, either sequentially or concurrently.
- 28 • *namespace* refers to a character string value assigned by the RM to a *job*. All *applications*
29 executed as part of that *job* share the same *namespace*. The *namespace* assigned to each *job* must
30 be unique within the scope of the governing RM and often is implemented as a string
31 representation of the numerical *emphJob ID*. The *namespace* and *job* terms will be used
32 interchangeably throughout the document.

- 1 • *application* represents a set of identical, but not necessarily unique, execution contexts within a
2 *job*.
- 3 • *process* is assumed for ease of presentation to be an operating system process, also commonly
4 referred to as a *heavyweight* process. A process is often comprised of multiple *lightweight*
5 *threads*, commonly known as simply *threads*. However, it is not the intent of the PMIx Standard
6 to restrict the term process to a particular concept or implementation.
- 7 • *client* refers to a process that was registered with the PMIx server prior to being started, and
8 connects to that PMIx server via `PMIx_Init` using its assigned namespace and rank with the
9 information required to connect to that server being provided to the process at time of start of
10 execution.
- 11 • *tool* refers to a process that may or may not have been registered with the PMIx server prior to
12 being started and initializes using `PMIx_tool_init`.
- 13 • *clone* refers to a process that was directly started by a PMIx client (e.g., using *fork/exec*) and calls
14 `PMIx_Init`, thus connecting to its local PMIx server using the same namespace and rank as its
15 parent process.
- 16 • *rank* refers to the numerical location (starting from zero) of a process within the defined scope.
17 Thus, *job rank* is the rank of a process within its *job* and is synonymous with its unqualified
18 *rank*, while *application rank* is the rank of that process within its *application*.
- 19 • *peer* refers to another process within the same *job*.
- 20 • *workflow* refers to an orchestrated execution plan typically involving multiple *jobs* carried out
21 under the control of a *workflow manager*. An example workflow might first execute a
22 computational job to generate the flow of liquid through a complex cavity, followed by a
23 visualization job that takes the output of the first job as its input to produce an image output.
- 24 • *scheduler* refers to the component of the SMS responsible for scheduling of resource allocations.
25 This is also generally referred to as the *system workflow manager* - for the purposes of this
26 document, the *WLM* acronym will be used interchangeably to refer to the scheduler.
- 27 • *resource manager* is used in a generic sense to represent the subsystem that will host the PMIx
28 server library. This could be a vendor-supplied resource manager or a third-party agent such as a
29 programming model's runtime library.
- 30 • *host environment* is used interchangeably with *resource manager* to refer to the process hosting
31 the PMIx server library.
- 32 • *node* refers to a single operating system instance. Note that this may encompass one or more
33 physical objects.
- 34 • *package* refers to a single object that is either soldered or connected to a printed circuit board via
35 a mechanical socket. Packages may contain multiple chips that include (but are not limited to)
36 processing units, memory, and peripheral interfaces.

- 1 • *processing unit*, or *PU*, is the electronic circuitry within a computer that executes instructions.
2 Depending upon architecture and configuration settings, it may consist of a single hardware
3 thread or multiple hardware threads collectively organized as a *core*.
 - 4 • *fabric* is used in a generic sense to refer to the networks within the system regardless of speed or
5 protocol. Any use of the term *network* in the document should be considered interchangeable
6 with *fabric*.
 - 7 • *fabric device* (or *fabric devices*) refers to an operating system fabric interface, which may be
8 physical or virtual. Any use of the term Network Interface Card (NIC) in the document should be
9 considered interchangeable with *fabric device*.
 - 10 • *fabric plane* refers to a collection of fabric devices in a common logical or physical
11 configuration. Fabric planes are often implemented in HPC clusters as separate overlay or
12 physical networks controlled by a dedicated fabric manager.
 - 13 • *attribute* refers to a key-value pair comprised of a string key (represented by a `pmix_key_t`
14 structure) and an associated value containing a PMIx data type (e.g., boolean, integer, or a more
15 complex PMIx structure). Attributes are used both as directives when passed as qualifiers to
16 APIs (e.g., in a `pmix_info_t` array), and to identify the contents of information (e.g., to
17 specify that the contents of the corresponding `pmix_value_t` in a `pmix_info_t` represent
18 the `PMIX_UNIV_SIZE`).
 - 19 • *key* refers to the string component of a defined *attribute*. The PMIx Standard will often refer to
20 passing of a *key* to an API (e.g., to the `PMIx_Query_info` or `PMIx_Get` APIs) as a means of
21 identifying requested information. In this context, the *data type* specified in the *attribute's*
22 definition indicates the data type the caller should expect to receive in return. Note that not all
23 *attributes* can be used as *keys* as some have specific uses solely as API qualifiers.
 - 24 • *instant on* refers to a PMIx concept defined as: "All information required for setup and
25 communication (including the address vector of endpoints for every process) is available to each
26 process at start of execution"
- 27 The following sections provide an overview of the conventions used throughout the PMIx Standard
28 document.

29 2.1 Notational Conventions

30 Some sections of this document describe programming language specific examples or APIs. Text
31 that applies only to programs for which the base language is C is shown as follows:

▼ C ▼

1 C specific text...

2 `int foo = 42;`

▲ C ▲

3 Some text is for information only, and is not part of the normative specification. These take several
4 forms, described in their examples below:

▼

5 Note: General text...

▲

▼ Rationale ▼

6 Throughout this document, the rationale for the design choices made in the interface specification is
7 set off in this section. Some readers may wish to skip these sections, while readers interested in
8 interface design may want to read them carefully.

▲

▼ Advice to users ▼

9 Throughout this document, material aimed at users and that illustrates usage is set off in this
10 section. Some readers may wish to skip these sections, while readers interested in programming
11 with the PMIx API may want to read them carefully.

▲

▼ Advice to PMIx library implementers ▼

12 Throughout this document, material that is primarily commentary to PMIx library implementers is
13 set off in this section. Some readers may wish to skip these sections, while readers interested in
14 PMIx implementations may want to read them carefully.

▲

▼ Advice to PMIx server hosts ▼

15 Throughout this document, material that is primarily commentary aimed at host environments (e.g.,
16 RMs and RunTime Environments (RTEs)) providing support for the PMIx server library is set off in
17 this section. Some readers may wish to skip these sections, while readers interested in integrating
18 PMIx servers into their environment may want to read them carefully.

▲

19 Attributes added in this version of the standard are shown in *magenta* to distinguish them from
20 those defined in prior versions, which are shown in *black*. Deprecated attributes are shown in *green*
21 and may be removed in a future version of the standard.

1 2.2 Semantics

2 The following terms will be taken to mean:

- 3 • *shall*, *must* and *will* indicate that the specified behavior is *required* of all conforming
4 implementations
- 5 • *should* and *may* indicate behaviors that a complete implementation would include, but are not
6 required of all conforming implementations

7 2.3 Naming Conventions

8 The PMIx standard has adopted the following conventions:

- 9 • PMIx constants and attributes are prefixed with "**PMIX_**".
- 10 • Structures and type definitions are prefixed with "**pmix_**".
- 11 • The string representation of attributes are prefixed with "**pmix**".
- 12 • Underscores are used to separate words in a function or variable name.
- 13 • Lowercase letters are used in PMIx client APIs except for the PMIx prefix (noted below) and the
14 first letter of the word following it. For example, **PMIx_Get_version**.
- 15 • PMIx server and tool APIs are all lower case letters following the prefix - e.g.,
16 **PMIx_server_register_nspace**.
- 17 • The **PMIX_** prefix is used to denote functions.
- 18 • The **pmix_** prefix is used to denote function pointer and type definitions.

19 Users shall not use the "**PMIX_**", "**PMIx_**", or "**pmix_**" prefixes for symbols in their code so as
20 to avoid symbol conflicts with PMIx implementations.

21 2.4 Procedure Conventions

22 While the current APIs are based on the C programming language, it is not the intent of the PMIx
23 Standard to preclude the use of other languages. Accordingly, the procedure specifications in the
24 PMIx Standard are written in a language-independent syntax with the arguments marked as IN,
25 OUT, or INOUT. The meanings of these are:

- 26 • IN: The call may use the input value but does not update the argument from the perspective of
27 the caller at any time during the calls execution,
- 28 • OUT: The call may update the argument but does not use its input value
- 29 • INOUT: The call may both use and update the argument.

1 Many PMIx interfaces, particularly nonblocking interfaces, use a **(void*)** callback data object
2 passed to the function that is then passed to the associated callback. On the client side, the callback
3 data object is an opaque, client-provided context that the client can pass to a non-blocking call.
4 When the nonblocking call completes, the callback data object is passed back to the client without
5 modification by the PMIx library, thus allowing the client to associate a context with that callback.
6 This is useful if there are many outstanding nonblocking calls.

7 A similar model is used for the server module functions (see [17.3.1](#)). In this case, the PMIx library
8 is making an upcall into its host via the PMIx server module callback function and passing a
9 specific callback function pointer and callback data object. The PMIx library expects the host to
10 call the cbfunc with the necessary arguments and pass back the original callback data object upon
11 completing the operation. This gives the server-side PMIx library the ability to associate a context
12 with the call back (since multiple operations may be outstanding). The host has no visibility into
13 the contents of the callback data object object, nor is permitted to alter it in any way.

CHAPTER 3

Data Structures and Types

1 This chapter defines PMIx standard data structures (along with macros for convenient use), types,
2 and constants. These apply to all consumers of the PMIx interface. Where necessary for
3 clarification, the description of, for example, an attribute may be copied from this chapter into a
4 section where it is used.

5 A PMIx implementation may define additional attributes beyond those specified in this document.

▼ Advice to PMIx library implementers ▼

6 Structures, types, and macros in the PMIx Standard are defined in terms of the C-programming
7 language. Implementers wishing to support other languages should provide the equivalent
8 definitions in a language-appropriate manner.

9 If a PMIx implementation chooses to define additional attributes they should avoid using the
10 "**PMIX**" prefix in their name or starting the attribute string with a "**pmix**" prefix. This helps the
11 end user distinguish between what is defined by the PMIx standard and what is specific to that
12 PMIx implementation, and avoids potential conflicts with attributes defined by the Standard.

▼ Advice to users ▼

13 Use of increment/decrement operations on indices inside PMIx macros is discouraged due to
14 unpredictable behavior as the index may be cited more than once in the macro. The PMIx standard
15 only governs the existence and syntax of macros - it does not specify their implementation.

16 Users are also advised to use the macros and APIs for creating, loading, and releasing PMIx
17 structures to avoid potential issues with release of memory. For example, pointing a
18 `pmix_envar_t` element at a static string variable and then using `PMIX_ENVAR_DESTRUCT` to
19 clear it would generate an error as the static string had not been allocated.

1 3.1 Constants

2 PMIx defines a few values that are used throughout the standard to set the size of fixed arrays or as
3 a means of identifying values with special meaning. The community makes every attempt to
4 minimize the number of such definitions. The constants defined in this section may be used before
5 calling any PMIx library initialization routine. Additional constants associated with specific data
6 structures or types are defined in the section describing that data structure or type.

7 **PMIX_MAX_NSLEN** Maximum namespace string length as an integer.

▼ **Advice to PMIx library implementers** ▼

8 **PMIX_MAX_NSLEN** should have a minimum value of 63 characters. Namespace arrays in PMIx
9 defined structures must reserve a space of size **PMIX_MAX_NSLEN**+1 to allow room for the **NULL**
10 terminator

11 **PMIX_MAX_KEYLEN** Maximum key string length as an integer.

▼ **Advice to PMIx library implementers** ▼

12 **PMIX_MAX_KEYLEN** should have a minimum value of 63 characters. Key arrays in PMIx defined
13 structures must reserve a space of size **PMIX_MAX_KEYLEN**+1 to allow room for the **NULL**
14 terminator

15 **PMIX_APP_WILDCARD** **UINT32_MAX** A value to indicate that the user wants the data for
16 the given key from every application that posted that key, or that the given value applies to all
17 applications within the given namespace.

1 3.1.1 PMIx Return Status Constants

2 The `pmix_status_t` type is an `int` compatible value for return status values. PMIx return
3 values other than `PMIX_SUCCESS` are required to always be negative. The return status value for a
4 successful operation is `PMIX_SUCCESS`, which must have an integer value of 0:

5 `PMIX_SUCCESS` 0 Success.

▼ Advice to PMIx library implementers ▼

6 A PMIx implementation must define all of the return status constants defined in the PMIx standard,
7 even if the implementation will never return the specific value to the caller.

▲

▼ Advice to users ▼

8 Other than `PMIX_SUCCESS` (which is required to be zero), the integer value of any PMIx error
9 constant is left to the PMIx library implementer with the constraint that it be negative and greater
10 magnitude (i.e. of larger absolute value) than `PMIX_EXTERNAL_ERR_BASE`. Thus, users are
11 advised to always refer to constants by name, and not by a specific implementation's integer value,
12 for portability between implementations and compatibility across library versions.

▲

13 The presentation of each API in this document includes a list of return status constants which are
14 either specific to that API or are expected to be returned by the API in normal use.

15 In addition, the following are general constants covering a variety of possible reasons an
16 implementation of an API may return a constant other than one of the constants presented with the
17 API. Although implementations can define and return additional error constants, implementations
18 are encouraged to return one of the return constants listed with the API or in the list presented here
19 to encourage portability across implementations.

20 `PMIX_ERROR` -1 General Error.

21 `PMIX_ERR_EXISTS` -11 The requested operation would overwrite an existing value -
22 typically returned when an operation would overwrite an existing file or directory.

23 `PMIX_ERR_EXISTS_OUTSIDE_SCOPE` -62 The requested key exists, but was posted in a
24 *scope* (see Section 5.2.1.1) that does not include the requester

25 `PMIX_ERR_INVALID_CRED` -12 Invalid security credentials.

26 `PMIX_ERR_WOULD_BLOCK` -15 Operation would block.

27 `PMIX_ERR_UNKNOWN_DATA_TYPE` -16 The data type specified in an input to the PMIx
28 library is not recognized by the implementation.

29 `PMIX_ERR_TYPE_MISMATCH` -18 The data type found in an object does not match the
30 expected data type as specified in the API call - e.g., a request to unpack a `PMIX_BOOL` value
31 from a buffer that does not contain a value of that type in the current unpack location.

32 `PMIX_ERR_UNPACK_INADEQUATE_SPACE` -19 Inadequate space to unpack data - the
33 number of values in the buffer exceeds the specified number to unpack.

1 **PMIX_ERR_UNPACK_READ_PAST_END_OF_BUFFER -50** Unpacking past the end of the
2 provided buffer - the number of values in the buffer is less than the specified number to
3 unpack, or a request was made to unpack a buffer beyond the buffer's end.
4 **PMIX_ERR_UNPACK_FAILURE -20** The unpack operation failed for an unspecified reason.
5 **PMIX_ERR_PACK_FAILURE -21** The pack operation failed for an unspecified reason.
6 **PMIX_ERR_NO_PERMISSIONS -23** The user lacks permissions to execute the specified
7 operation.
8 **PMIX_ERR_TIMEOUT -24** Either a user-specified or system-internal timeout expired.
9 **PMIX_ERR_UNREACH -25** The specified target server or client process is not reachable -
10 i.e., a suitable connection either has not been or can not be made.
11 **PMIX_ERR_BAD_PARAM -27** One or more incorrect parameters (e.g., passing an attribute
12 with a value of the wrong type), or multiple parameters containing conflicting directives (e.g.,
13 multiple instances of the same attribute with different values, or different attributes specifying
14 conflicting behaviors), were passed to a PMIx API.
15 **PMIX_ERR_EMPTY -60** An array or list was given that has no members in it - i.e., the object
16 is empty.
17 **PMIX_ERR_RESOURCE_BUSY -28** Resource busy - typically seen when an attempt to
18 establish a connection to another process (e.g., a PMIx server) cannot be made due to a
19 communication failure.
20 **PMIX_ERR_OUT_OF_RESOURCE -29** Resource exhausted.
21 **PMIX_ERR_INIT -31** The requested operation requires that the PMIx library be initialized
22 prior to being called.
23 **PMIX_ERR_NOMEM -32** Out of memory.
24 **PMIX_ERR_NOT_FOUND -46** The requested information was not found.
25 **PMIX_ERR_NOT_SUPPORTED -47** The requested operation is not supported by either the
26 PMIx implementation or the host environment.
27 **PMIX_ERR_PARAM_VALUE_NOT_SUPPORTED -59** The requested operation is supported
28 by the PMIx implementation and (if applicable) the host environment. However, at least one
29 supplied parameter was given an unsupported value, and the operation cannot therefore be
30 executed as requested.
31 **PMIX_ERR_COMM_FAILURE -49** Communication failure - a message failed to be sent or
32 received, but the connection remains intact.
33 **PMIX_ERR_LOST_CONNECTION -61** Lost connection between server and client or tool.
34 **PMIX_ERR_INVALID_OPERATION -158** The requested operation is supported by the
35 implementation and host environment, but fails to meet a requirement (e.g., requesting to
36 *disconnect* from processes without first *connecting* to them, inclusion of conflicting
37 directives, or a request to perform an operation that conflicts with an ongoing one).
38 **PMIX_OPERATION_IN_PROGRESS -156** A requested operation is already in progress -
39 the duplicate request shall therefore be ignored.
40 **PMIX_OPERATION_SUCCEEDED -157** The requested operation was performed atomically
41 - no callback function will be executed.

1 **PMIX_ERR_PARTIAL_SUCCESS** -52 The operation is considered successful but not all
2 elements of the operation were concluded (e.g., some members of a group construct operation
3 chose not to participate).

3.1.1.1 User-Defined Error and Event Constants

5 PMIx establishes a boundary for constants defined in the PMIx standard. Negative values larger
6 (i.e., more negative) than this (and any positive values greater than zero) are guaranteed not to
7 conflict with PMIx values.

8 **PMIX_EXTERNAL_ERR_BASE** A starting point for user-level defined error and event
9 constants. Negative values that are more negative than the defined constant are guaranteed not
10 to conflict with PMIx values. Definitions should always be based on the
11 **PMIX_EXTERNAL_ERR_BASE** constant and not a specific value as the value of the constant
12 may change.

3.2 Data Types

14 This section defines various data types used by the PMIx APIs. The version of the standard in
15 which a particular data type was introduced is shown in the margin.

3.2.1 Key Structure

17 The `pmix_key_t` structure is a statically defined character array of length
18 **PMIX_MAX_KEYLEN**+1, thus supporting keys of maximum length **PMIX_MAX_KEYLEN** while
19 preserving space for a mandatory **NULL** terminator.

PMIx v2.0

```
typedef char pmix_key_t[PMIX_MAX_KEYLEN+1];
```

21 Characters in the key must be standard alphanumeric values supported by common utilities such as
22 *strcmp*.

Advice to users

23 References to keys in PMIx v1 were defined simply as an array of characters of size
24 **PMIX_MAX_KEYLEN**+1. The `pmix_key_t` type definition was introduced in version 2 of the
25 standard. The two definitions are code-compatible and thus do not represent a break in backward
26 compatibility.

27 Passing a `pmix_key_t` value to the standard *sizeof* utility can result in compiler warnings of
28 incorrect returned value. Users are advised to avoid using *sizeof(pmix_key_t)* and instead rely on
29 the **PMIX_MAX_KEYLEN** constant.

1 3.2.1.1 Key support macros

2 The following macros are provided for convenience when working with PMIx keys.

3 Check key macro

4 Compare the key in a `pmix_info_t` to a given value.

PMIx v3.0

▼ `C` _____ ▼

5 `PMIX_CHECK_KEY(a, b)`

▲ _____ `C` _____ ▲

6 **IN a**

7 Pointer to the structure whose key is to be checked (pointer to `pmix_info_t`)

8 **IN b**

9 String value to be compared against (`char*`)

10 Returns `true` if the key matches the given value

11 Check reserved key macro

12 Check if the given key is a PMIx *reserved* key as described in Chapter 6.

PMIx v4.0

▼ `C` _____ ▼

13 `PMIX_CHECK_RESERVED_KEY(a)`

▲ _____ `C` _____ ▲

14 **IN a**

15 String value to be checked (`char*`)

16 Returns `true` if the key is reserved by the Standard.

17 Load key macro

18 Load a key into a `pmix_info_t`.

PMIx v4.0

▼ `C` _____ ▼

19 `PMIX_LOAD_KEY(a, b)`

▲ _____ `C` _____ ▲

20 **IN a**

21 Pointer to the structure whose key is to be loaded (pointer to `pmix_info_t`)

22 **IN b**

23 String value to be loaded (`char*`)

24 No return value.

1 3.2.2 Namespace Structure

2 The `pmix_nspace_t` structure is a statically defined character array of length
3 `PMIX_MAX_NSLEN+1`, thus supporting namespaces of maximum length `PMIX_MAX_NSLEN`
4 while preserving space for a mandatory `NULL` terminator.

```
▼ C  
5 typedef char pmix_nspace_t[PMIX_MAX_NSLEN+1];  
▲ C
```

6 Characters in the namespace must be standard alphanumeric values supported by common utilities
7 such as `strcmp`.

▼ Advice to users ▲

8 References to namespace values in PMIx v1 were defined simply as an array of characters of size
9 `PMIX_MAX_NSLEN+1`. The `pmix_nspace_t` type definition was introduced in version 2 of the
10 standard. The two definitions are code-compatible and thus do not represent a break in backward
11 compatibility.

12 Passing a `pmix_nspace_t` value to the standard `sizeof` utility can result in compiler warnings of
13 incorrect returned value. Users are advised to avoid using `sizeof(pmix_nspace_t)` and instead rely
14 on the `PMIX_MAX_NSLEN` constant.

15 3.2.2.1 Namespace support macros

16 The following macros are provided for convenience when working with PMIx namespace
17 structures.

18 Check namespace macro

19 Compare the string in a `pmix_nspace_t` to a given value.

PMIx v3.0

```
▼ C  
20 PMIX_CHECK_NAMESPACE(a, b)  
▲ C
```

21 **IN a**
22 Pointer to the structure whose value is to be checked (pointer to `pmix_nspace_t`)

23 **IN b**
24 String value to be compared against (`char*`)

25 Returns `true` if the namespace matches the given value

1 **Check invalid namespace macro**
2 Check if the provided `pmix_namespace_t` is invalid.

▼ C ▼

3 **PMIX_NAMESPACE_INVALID (a)**

▲ C ▲

4 **IN a**
5 Pointer to the structure whose value is to be checked (pointer to `pmix_namespace_t`)

6 Returns `true` if the namespace is invalid (i.e., starts with a `NULL` resulting in a zero-length string
7 value)

8 **Load namespace macro**
9 Load a namespace into a `pmix_namespace_t`.

PMIx v4.0 ▼ C ▼

10 **PMIX_LOAD_NAMESPACE (a, b)**

▲ C ▲

11 **IN a**
12 Pointer to the target structure (pointer to `pmix_namespace_t`)

13 **IN b**
14 String value to be loaded - if `NULL` is given, then the target structure will be initialized to
15 zero's (`char*`)

16 No return value.

17 3.2.3 Rank Structure

18 The `pmix_rank_t` structure is a `uint32_t` type for rank values.

PMIx v1.0 ▼ C ▼

19 **typedef uint32_t pmix_rank_t;**

▲ C ▲

20 The following constants can be used to set a variable of the type `pmix_rank_t`. All definitions
21 were introduced in version 1 of the standard unless otherwise marked. Valid rank values start at
22 zero.

23 **PMIX_RANK_UNDEF UINT32_MAX** A value to request job-level data where the information
24 itself is not associated with any specific rank, or when passing a `pmix_proc_t` identifier to
25 an operation that only references the namespace field of that structure.

26 **PMIX_RANK_WILDCARD UINT32_MAX-1** A value to indicate that the user wants the data
27 for the given key from every rank that posted that key.

28 **PMIX_RANK_LOCAL_NODE UINT32_MAX-2** Special rank value used to define groups of
29 ranks. This constant defines the group of all ranks on a local node.

1 **PMIX_RANK_LOCAL_PEERS** **UINT32_MAX-4** Special rank value used to define groups of
 2 ranks. This constant defines the group of all ranks on a local node within the same namespace
 3 as the current process.
 4 **PMIX_RANK_INVALID** **UINT32_MAX-3** An invalid rank value.
 5 **PMIX_RANK_VALID** **UINT32_MAX-50** Define an upper boundary for valid rank values.

6 3.2.3.1 Rank support macros

7 The following macros are provided for convenience when working with PMIx ranks.

8 **Check rank macro**

9 Check two ranks for equality, taking into account wildcard values

PMIx v4.0  **C** 

10 **PMIX_CHECK_RANK**(a, b)

 **C** 

11 **IN** a

12 Rank to be checked (**pmix_rank_t**)

13 **IN** b

14 Rank to be checked (**pmix_rank_t**)

15 Returns **true** if the ranks are equal, or at least one of the ranks is **PMIX_RANK_WILDCARD**

16 **Check rank is valid macro**

17 Check if the given rank is a valid value

Provisional  **C** 
v4.1

18 **PMIX_RANK_IS_VALID**(a)

 **C** 

19 **IN** a

20 Rank to be checked (**pmix_rank_t**)

21 Returns **true** if the given rank is valid (i.e., less than **PMIX_RANK_VALID**)

22 3.2.4 Process Structure

23 The **pmix_proc_t** structure is used to identify a single process in the PMIx universe. It contains
 24 a reference to the namespace and the **pmix_rank_t** within that namespace.

PMIx v1.0  **C** 

```
25 typedef struct pmix_proc {
26     pmix_namespace_t nspace;
27     pmix_rank_t rank;
28 } pmix_proc_t;
```

 **C** 

1 3.2.4.1 Process structure support macros

2 The following macros are provided to support the `pmix_proc_t` structure.

3 **Static initializer for the proc structure**

4 Provide a static initializer for the `pmix_proc_t` fields.

PMIx v4.2

▼ C

5 **PMIX_PROC_STATIC_INIT**

▲ C

6 **Initialize the proc structure**

7 Initialize the `pmix_proc_t` fields.

PMIx v1.0

▼ C

8 **PMIX_PROC_CONSTRUCT (m)**

▲ C

9 **IN** m

10 Pointer to the structure to be initialized (pointer to `pmix_proc_t`)

11 **Destruct the proc structure**

12 Destruct the `pmix_proc_t` fields.

▼ C

13 **PMIX_PROC_DESTRUCT (m)**

▲ C

14 **IN** m

15 Pointer to the structure to be destructed (pointer to `pmix_proc_t`)

16 There is nothing to release here as the fields in `pmix_proc_t` are either a statically-declared array
17 (the namespace) or a single value (the rank). However, the macro is provided for symmetry in the
18 code and for future-proofing should some allocated field be included some day.

19 **Create a proc array**

20 Allocate and initialize an array of `pmix_proc_t` structures.

PMIx v1.0

▼ C

21 **PMIX_PROC_CREATE (m, n)**

▲ C

22 **INOUT** m

23 Address where the pointer to the array of `pmix_proc_t` structures shall be stored (handle)

24 **IN** n

25 Number of structures to be allocated (`size_t`)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22

Free a proc structure

Release a `pmix_proc_t` structure.



PMIX_PROC_RELEASE (m)



IN m

Pointer to a `pmix_proc_t` structure (handle)

Free a proc array

Release an array of `pmix_proc_t` structures.

PMIx v1.0



PMIX_PROC_FREE (m, n)



IN m

Pointer to the array of `pmix_proc_t` structures (handle)

IN n

Number of structures in the array (`size_t`)

Load a proc structure

Load values into a `pmix_proc_t`.

PMIx v2.0



PMIX_PROC_LOAD (m, n, r)



IN m

Pointer to the structure to be loaded (pointer to `pmix_proc_t`)

IN n

Namespace to be loaded (`pmix_namespace_t`)

IN r

Rank to be assigned (`pmix_rank_t`)

No return value. Deprecated in favor of `PMIX_LOAD_PROCID`

1 **Compare identifiers**
2 Compare two `pmix_proc_t` identifiers.

▼ C ————— ▼

3 **PMIX_CHECK_PROCID (a, b)**

▲ C ————— ▲

4 **IN a**
5 Pointer to a structure whose ID is to be compared (pointer to `pmix_proc_t`)
6 **IN b**
7 Pointer to a structure whose ID is to be compared (pointer to `pmix_proc_t`)

8 Returns `true` if the two structures contain matching namespaces and:

- 9 • the ranks are the same value
- 10 • one of the ranks is `PMIX_RANK_WILDCARD`

11 **Check if a process identifier is valid**

12 Check for invalid namespace or rank value

Provisional
v4.1

▼ C ————— ▼

13 **PMIX_PROCID_INVALID (a)**

▲ C ————— ▲

14 **IN a**
15 Pointer to a structure whose ID is to be checked (pointer to `pmix_proc_t`)

16 Returns `true` if the process identifier contains either an empty (i.e., invalid) `namespace` field or a `rank`
17 field of `PMIX_RANK_INVALID`

18 **Load a proclD structure**

19 Load values into a `pmix_proc_t`.

PMIx v4.0

▼ C ————— ▼

20 **PMIX_LOAD_PROCID (m, n, r)**

▲ C ————— ▲

21 **IN m**
22 Pointer to the structure to be loaded (pointer to `pmix_proc_t`)
23 **IN n**
24 Namespace to be loaded (`pmix_namespace_t`)
25 **IN r**
26 Rank to be assigned (`pmix_rank_t`)

1 **Transfer a proclD structure**

2 Transfer contents of one `pmix_proc_t` value to another `pmix_proc_t`.

▼ `PMIX_PROCID_XFER(d, s)` C

3 **PMIX_PROCID_XFER(d, s)**

▲ `PMIX_PROCID_XFER(d, s)` C

4 **IN d**

5 Pointer to the target structure (pointer to `pmix_proc_t`)

6 **IN s**

7 Pointer to the source structure (pointer to `pmix_proc_t`)

8 **Construct a multi-cluster namespace**

9 Construct a multi-cluster identifier containing a cluster ID and a namespace.

PMIx v4.0 ▼ `PMIX_MULTICLUSTER_NAMESPACE_CONSTRUCT(m, n, r)` C

10 **PMIX_MULTICLUSTER_NAMESPACE_CONSTRUCT(m, n, r)**

▲ `PMIX_MULTICLUSTER_NAMESPACE_CONSTRUCT(m, n, r)` C

11 **IN m**

12 `pmix_namespace_t` structure that will contain the multi-cluster identifier (`pmix_namespace_t`)

13 **IN n**

14 Cluster identifier (`char*`)

15 **IN r**

16 Namespace to be loaded (`pmix_namespace_t`)

17 Combined length of the cluster identifier and namespace must be less than `PMIX_MAX_NSLEN-2`.

18 **Parse a multi-cluster namespace**

19 Parse a multi-cluster identifier into its cluster ID and namespace parts.

PMIx v4.0 ▼ `PMIX_MULTICLUSTER_NAMESPACE_PARSE(m, n, r)` C

20 **PMIX_MULTICLUSTER_NAMESPACE_PARSE(m, n, r)**

▲ `PMIX_MULTICLUSTER_NAMESPACE_PARSE(m, n, r)` C

21 **IN m**

22 `pmix_namespace_t` structure containing the multi-cluster identifier (pointer to `pmix_namespace_t`)

23 **IN n**

24 Location where the cluster ID is to be stored (`pmix_namespace_t`)

25 **IN r**

26 Location where the namespace is to be stored (`pmix_namespace_t`)

27

1 3.2.5 Process State Structure

2 The `pmix_proc_state_t` structure is a `uint8_t` type for process state values. The following
3 constants can be used to set a variable of the type `pmix_proc_state_t`.

Advice to users

4 The fine-grained nature of the following constants may exceed the ability of an RM to provide
5 updated process state values during the process lifetime. This is particularly true of states for
6 short-lived processes.

7 `PMIX_PROC_STATE_UNDEF 0` Undefined process state.
8 `PMIX_PROC_STATE_PREPPED 1` Process is ready to be launched.
9 `PMIX_PROC_STATE_LAUNCH_UNDERWAY 2` Process launch is underway.
10 `PMIX_PROC_STATE_RESTART 3` Process is ready for restart.
11 `PMIX_PROC_STATE_TERMINATE 4` Process is marked for termination.
12 `PMIX_PROC_STATE_RUNNING 5` Process has been locally `fork`'ed by the RM.
13 `PMIX_PROC_STATE_CONNECTED 6` Process has connected to PMIx server.
14 `PMIX_PROC_STATE_UNTERMINATED` Define a "boundary" between the terminated states
15 and `PMIX_PROC_STATE_CONNECTED` so users can easily and quickly determine if a
16 process is still running or not. Any value less than this constant means that the process has not
17 terminated.
18 `PMIX_PROC_STATE_TERMINATED` Process has terminated and is no longer running.
19 `PMIX_PROC_STATE_ERROR` Define a boundary so users can easily and quickly determine if
20 a process abnormally terminated. Any value above this constant means that the process has
21 terminated abnormally.
22 `PMIX_PROC_STATE_KILLED_BY_CMD 51` Process was killed by a command.
23 `PMIX_PROC_STATE_ABORTED 52` Process was aborted by a call to `PMIx_Abort`.
24 `PMIX_PROC_STATE_FAILED_TO_START 53` Process failed to start.
25 `PMIX_PROC_STATE_ABORTED_BY_SIG 54` Process aborted by a signal.
26 `PMIX_PROC_STATE_TERM_WO_SYNC 55` Process exited without calling
27 `PMIx_Finalize`.
28 `PMIX_PROC_STATE_COMM_FAILED 56` Process communication has failed.
29 `PMIX_PROC_STATE_SENSOR_BOUND_EXCEEDED 57` Process exceeded a specified
30 sensor limit.
31 `PMIX_PROC_STATE_CALLED_ABORT 58` Process called `PMIx_Abort`.
32 `PMIX_PROC_STATE_HEARTBEAT_FAILED 59` Process failed to send heartbeat within
33 specified time limit.
34 `PMIX_PROC_STATE_MIGRATING 60` Process failed and is waiting for resources before
35 restarting.
36 `PMIX_PROC_STATE_CANNOT_RESTART 61` Process failed and cannot be restarted.
37 `PMIX_PROC_STATE_TERM_NON_ZERO 62` Process exited with a non-zero status.
38 `PMIX_PROC_STATE_FAILED_TO_LAUNCH 63` Unable to launch process.

1 3.2.6 Process Information Structure

2 The `pmix_proc_info_t` structure defines a set of information about a specific process
3 including its name, location, and state.

PMIx v2.0

```
4 typedef struct pmix_proc_info {  
5     /** Process structure */  
6     pmix_proc_t proc;  
7     /** Hostname where process resides */  
8     char *hostname;  
9     /** Name of the executable */  
10    char *executable_name;  
11    /** Process ID on the host */  
12    pid_t pid;  
13    /** Exit code of the process. Default: 0 */  
14    int exit_code;  
15    /** Current state of the process */  
16    pmix_proc_state_t state;  
17 } pmix_proc_info_t;
```

18 3.2.6.1 Process information structure support macros

19 The following macros are provided to support the `pmix_proc_info_t` structure.

20 **Static initializer for the proc info structure**

21 *Provisional*

Provide a static initializer for the `pmix_proc_info_t` fields.

PMIx v4.2

```
22 PMIX_PROC_INFO_STATIC_INIT
```

23 **Initialize the process information structure**

24 Initialize the `pmix_proc_info_t` fields.

PMIx v2.0

```
25 PMIX_PROC_INFO_CONSTRUCT (m)
```

26 **IN** m

27 Pointer to the structure to be initialized (pointer to `pmix_proc_info_t`)

1 **Destruct the process information structure**

2 Destruct the `pmix_proc_info_t` fields.



3 **PMIX_PROC_INFO_DESTRUCT (m)**



4 **IN m**

5 Pointer to the structure to be destructed (pointer to `pmix_proc_info_t`)

6 **Create a process information array**

7 Allocate and initialize a `pmix_proc_info_t` array.

PMIx v2.0



8 **PMIX_PROC_INFO_CREATE (m, n)**



9 **INOUT m**

10 Address where the pointer to the array of `pmix_proc_info_t` structures shall be stored
11 (handle)

12 **IN n**

13 Number of structures to be allocated (`size_t`)

14 **Free a process information structure**

15 Release a `pmix_proc_info_t` structure.

PMIx v2.0



16 **PMIX_PROC_INFO_RELEASE (m)**



17 **IN m**

18 Pointer to a `pmix_proc_info_t` structure (handle)

19 **Free a process information array**

20 Release an array of `pmix_proc_info_t` structures.

PMIx v2.0



21 **PMIX_PROC_INFO_FREE (m, n)**



22 **IN m**

23 Pointer to the array of `pmix_proc_info_t` structures (handle)

24 **IN n**

25 Number of structures in the array (`size_t`)

1 3.2.7 Job State Structure

2 The `pmix_job_state_t` structure is a `uint8_t` type for job state values. The following
3 constants can be used to set a variable of the type `pmix_job_state_t`.

Advice to users

4 The fine-grained nature of the following constants may exceed the ability of an RM to provide
5 updated job state values during the job lifetime. This is particularly true for short-lived jobs.

6 **PMIX_JOB_STATE_UNDEF** 0 Undefined job state.
7 **PMIX_JOB_STATE_AWAITING_ALLOC** 1 Job is waiting for resources to be allocated to it.
8 **PMIX_JOB_STATE_LAUNCH_UNDERWAY** 2 Job launch is underway.
9 **PMIX_JOB_STATE_RUNNING** 3 All processes in the job have been spawned and are
10 executing.
11 **PMIX_JOB_STATE_SUSPENDED** 4 All processes in the job have been suspended.
12 **PMIX_JOB_STATE_CONNECTED** 5 All processes in the job have connected to their PMIx
13 server.
14 **PMIX_JOB_STATE_UNTERMINATED** Define a “boundary” between the terminated states
15 and **PMIX_JOB_STATE_TERMINATED** so users can easily and quickly determine if a job is
16 still running or not. Any value less than this constant means that the job has not terminated.
17 **PMIX_JOB_STATE_TERMINATED** All processes in the job have terminated and are no
18 longer running - typically will be accompanied by the job exit status in response to a query.
19 **PMIX_JOB_STATE_TERMINATED_WITH_ERROR** Define a boundary so users can easily
20 and quickly determine if a job abnormally terminated - typically will be accompanied by a
21 job-related error code in response to a query Any value above this constant means that the job
22 terminated abnormally.

23 3.2.8 Value Structure

24 The `pmix_value_t` structure is used to represent the value passed to **PMIx_Put** and retrieved
25 by **PMIx_Get**, as well as many of the other PMIx functions.

26 A collection of values may be specified under a single key by passing a `pmix_value_t`
27 containing an array of type `pmix_data_array_t`, with each array element containing its own
28 object. All members shown below were introduced in version 1 of the standard unless otherwise
29 marked.

PMIx v1.0

```

1  typedef struct pmix_value {
2      pmix_data_type_t type;
3      union {
4          bool flag;
5          uint8_t byte;
6          char *string;
7          size_t size;
8          pid_t pid;
9          int integer;
10         int8_t int8;
11         int16_t int16;
12         int32_t int32;
13         int64_t int64;
14         unsigned int uint;
15         uint8_t uint8;
16         uint16_t uint16;
17         uint32_t uint32;
18         uint64_t uint64;
19         float fval;
20         double dval;
21         struct timeval tv;
22         time_t time; // version 2.0
23         pmix_status_t status; // version 2.0
24         pmix_rank_t rank; // version 2.0
25         pmix_proc_t *proc; // version 2.0
26         pmix_byte_object_t bo;
27         pmix_persistence_t persist; // version 2.0
28         pmix_scope_t scope; // version 2.0
29         pmix_data_range_t range; // version 2.0
30         pmix_proc_state_t state; // version 2.0
31         pmix_proc_info_t *pinfo; // version 2.0
32         pmix_data_array_t *darray; // version 2.0
33         void *ptr; // version 2.0
34         pmix_alloc_directive_t adir; // version 2.0
35     } data;
36 } pmix_value_t;

```

37 3.2.8.1 Value structure support

38 The following macros and APIs are provided to support the `pmix_value_t` structure.

1 **Static initializer for the value structure**

2 Provide a static initializer for the `pmix_value_t` fields.



3 `PMIX_VALUE_STATIC_INIT`



4 **Initialize the value structure**

5 Initialize the `pmix_value_t` fields.

PMIx v1.0



6 `PMIX_VALUE_CONSTRUCT (m)`



7 **IN** `m`

8 Pointer to the structure to be initialized (pointer to `pmix_value_t`)

9 **Destruct the value structure**

10 Destruct the `pmix_value_t` fields.

PMIx v1.0



11 `PMIX_VALUE_DESTRUCT (m)`



12 **IN** `m`

13 Pointer to the structure to be destructed (pointer to `pmix_value_t`)

14 **Create a value array**

15 Allocate and initialize an array of `pmix_value_t` structures.

PMIx v1.0



16 `PMIX_VALUE_CREATE (m, n)`



17 **INOUT** `m`

18 Address where the pointer to the array of `pmix_value_t` structures shall be stored (handle)

19 **IN** `n`

20 Number of structures to be allocated (`size_t`)

21 **Free a value structure**

22 Release a `pmix_value_t` structure.

PMIx v4.0



23 `PMIX_VALUE_RELEASE (m)`



24 **IN** `m`

25 Pointer to a `pmix_value_t` structure (handle)

1 **Free a value array**
2 Release an array of `pmix_value_t` structures.

PMIx v1.0

▼ `PMIX_VALUE_FREE` C

3 **PMIX_VALUE_FREE**(`m`, `n`)

▲ `PMIX_VALUE_FREE` C

4 **IN** `m`
5 Pointer to the array of `pmix_value_t` structures (handle)
6 **IN** `n`
7 Number of structures in the array (`size_t`)

8 **Load a value structure**

9 **Summary**

10 Load data into a `pmix_value_t` structure.

11 **Format**

PMIx v4.2

▼ `PMIX_VALUE_LOAD` C

12 `pmix_status_t`
13 `PMIX_Value_load`(`pmix_value_t *val`,
14 `const void *data`,
15 `pmix_data_type_t type`);

▲ `PMIX_VALUE_LOAD` C

16 **IN** `val`
17 The `pmix_value_t` into which the data is to be loaded (pointer to `pmix_value_t`)
18 **IN** `data`
19 Pointer to the data value to be loaded (handle)
20 **IN** `type`
21 Type of the provided data value (`pmix_data_type_t`)

22 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

23 **Description**

24 Copy the provided data into the `pmix_value_t`. Any data stored in the source value can be
25 modified or free'd without affecting the copied data once the function has completed.

26 **Unload a value structure**

27 **Summary**

28 Unload data from a `pmix_value_t` structure.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29

Format

C

```
pmix_status_t  
PMIx_Value_unload(pmix_value_t *val,  
                  void **data,  
                  size_t *sz);
```

C

IN val

The `pmix_value_t` from which the data is to be unloaded (pointer to `pmix_value_t`)

INOUT data

Pointer to the location where the data value is to be returned (handle)

INOUT sz

Pointer to return the size of the unloaded value (handle)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

Return a copy of the data in the `pmix_value_t`. The source value can be modified or free'd without affecting the copied data once the function has completed.

Advice to users

Memory will be allocated and the pointer to that data will be in the `data` argument - the source `pmix_value_t` will not be altered. The user is responsible for releasing the returned data.

Transfer data between value structures

Summary

Transfer the data value between two `pmix_value_t` structures.

Format

C

PMIx v4.2

```
pmix_status_t  
PMIx_Value_xfer(pmix_value_t *dest,  
                const pmix_value_t *src);
```

C

IN dest

Pointer to the `pmix_value_t` destination (handle)

IN src

Pointer to the `pmix_value_t` source (handle)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17

18
19
20
21
22
23
24
25

26
27

Description

Copy the data in the source `pmix_value_t` into the destination `pmix_value_t`. The source value can be modified or free'd without affecting the copied data once the function has completed.

Retrieve a numerical value from a value struct

Retrieve a numerical value from a `pmix_value_t` structure.

PMIx v3.0

```
▼ _____ C _____ ▼  
PMIX_VALUE_GET_NUMBER(s, m, n, t)  
▲ _____ C _____ ▲
```

OUT `s`

Status code for the request (`pmix_status_t`)

IN `m`

Pointer to the `pmix_value_t` structure (handle)

OUT `n`

Variable to be set to the value (match expected type)

IN `t`

Type of number expected in `m` (`pmix_data_type_t`)

Sets the provided variable equal to the numerical value contained in the given `pmix_value_t`, returning success if the data type of the value matches the expected type and `PMIX_ERR_BAD_PARAM` if it doesn't

3.2.9 Info Structure

The `pmix_info_t` structure defines a key/value pair with associated directive. All fields were defined in version 1.0 unless otherwise marked.

PMIx v1.0

```
▼ _____ C _____ ▼  
typedef struct pmix_info_t {  
    pmix_key_t key;  
    pmix_info_directives_t flags; // version 2.0  
    pmix_value_t value;  
} pmix_info_t;  
▲ _____ C _____ ▲
```

3.2.9.1 Info structure support macros

The following macros are provided to support the `pmix_info_t` structure.

1 **Static initializer for the info structure**

2 Provide a static initializer for the `pmix_info_t` fields.



3 **PMIX_INFO_STATIC_INIT**



4 **Initialize the info structure**

5 Initialize the `pmix_info_t` fields.

PMIx v1.0



6 **PMIX_INFO_CONSTRUCT (m)**



7 **IN** m

8 Pointer to the structure to be initialized (pointer to `pmix_info_t`)

9 **Destruct the info structure**

10 Destruct the `pmix_info_t` fields.

PMIx v1.0



11 **PMIX_INFO_DESTRUCT (m)**



12 **IN** m

13 Pointer to the structure to be destructed (pointer to `pmix_info_t`)

14 **Create an info array**

15 Allocate and initialize an array of info structures.

PMIx v1.0



16 **PMIX_INFO_CREATE (m, n)**



17 **INOUT** m

18 Address where the pointer to the array of `pmix_info_t` structures shall be stored (handle)

19 **IN** n

20 Number of structures to be allocated (`size_t`)

1 **Free an info array**
2 Release an array of `pmix_info_t` structures.

▼ `PMIX_INFO_FREE` C

3 **PMIX_INFO_FREE**(`m`, `n`)

▲ `PMIX_INFO_FREE` C

4 **IN** `m`
5 Pointer to the array of `pmix_info_t` structures (handle)
6 **IN** `n`
7 Number of structures in the array (`size_t`)

8 **Load key and value data into a info struct**
9 **Summary**

10 Load data into a `pmix_info_t` structure.

11 **Format**

PMIx v4.2

▼ `PMIX_INFO_LOAD` C

12 `pmix_status_t`
13 `PMIx_Info_load`(`pmix_info_t *info`,
14 `const char* key`,
15 `const void *data`,
16 `pmix_data_type_t type`);

▲ `PMIX_INFO_LOAD` C

17 **IN** `info`
18 The `pmix_info_t` into which the data is to be loaded (handle)
19 **IN** `key`
20 Pointer to the key to be loaded (handle)
21 **IN** `data`
22 Pointer to the data value to be loaded (handle)
23 **IN** `type`
24 Type of the provided data value (`pmix_data_type_t`)

25 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

26 **Description**

27 Copy the provided data into the `pmix_info_t`. Any data stored in the source parameters can be
28 modified or free'd without affecting the copied data once the function has completed.

29 **Copy data between info structures**

30 **Summary**

31 Copy all data between two `pmix_info_t` structures.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

Format

C

```
pmix_status_t
PMIx_Info_xfer (pmix_info_t *dest,
                pmix_info_t *src);
```

C

IN dest

The `pmix_info_t` into which the data is to be copied (handle)

IN src

The `pmix_info_t` from which the data is to be copied (handle)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

Copy the data in the source `pmix_info_t` into the destination. Any data stored in the source structure can be modified or free'd without affecting the copied data once the function has completed.

Test a boolean info struct

A special macro for checking if a boolean `pmix_info_t` is `true`.

PMIx v2.0

C

PMIX_INFO_TRUE (m)

C

IN m

Pointer to a `pmix_info_t` structure (handle)

A `pmix_info_t` structure is considered to be of type `PMIX_BOOL` and value `true` if:

- the structure reports a type of `PMIX_UNDEF`, or
- the structure reports a type of `PMIX_BOOL` and the data flag is `true`

3.2.9.2 Info structure list macros

Constructing an array of `pmix_info_t` is a fairly common operation. The following macros are provided to simplify this construction.

Start a list of `pmix_info_t` structures

Summary

Initialize a list of `pmix_info_t` structures. The actual list is opaque to the caller and is implementation-dependent.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32

Format

C

```
void*
PMIx_Info_list_start(void);
```

C

Description

Note that the returned pointer will be initialized to an opaque structure whose elements are implementation-dependent. The caller must not modify or dereference the object.

Add a `pmix_info_t` structure to a list

Summary

Add a `pmix_info_t` structure containing the specified value to the provided list.

Format

C

```
pmix_status_t
PMIx_Info_list_add(void *ptr,
                   const char *key,
                   const void *value,
                   pmix_data_type_t type);
```

C

IN ptr

A `void*` pointer initialized via `PMIx_Info_list_start` (handle)

IN key

String key to be loaded - must be less than or equal to `PMIX_MAX_KEYLEN` in length (handle)

IN value

Pointer to the data value to be loaded (handle)

IN type

Type of the provided data value (`pmix_data_type_t`)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

Copy the provided key and data into a `pmix_info_t` on the list. The key and any data stored in the source value can be modified or free'd without affecting the copied data once the function has completed.

Transfer a `pmix_info_t` structure to a list

Summary

Transfer the information in a `pmix_info_t` structure to a structure on the provided list.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32

Format

C

```
pmix_status_t  
PMIx_Info_list_xfer(void *ptr,  
                    const pmix_info_t *src);
```

C

IN ptr

A void* pointer initialized via `PMIx_Info_list_start` (handle)

IN src

Pointer to the source `pmix_info_t` (pointer to `pmix_info_t`)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

All data (including key, value, and directives) will be copied into a destination `pmix_info_t` on the list. The source `pmix_info_t` may be free'd without affecting the copied data once the function has completed.

Convert a `pmix_info_t` list to an array

Summary

Transfer the information in the provided `pmix_info_t` list to a `pmix_data_array_t` array

Format

C

PMIx v4.2

```
pmix_status_t  
PMIx_Info_list_convert(void *ptr,  
                       pmix_data_array_t *par);
```

C

IN ptr

A void* pointer initialized via `PMIx_Info_list_start` (handle)

IN par

Pointer to an instantiated `pmix_data_array_t` structure where the `pmix_info_t` array is to be stored (pointer to `pmix_data_array_t`)

Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Description

Information collected in the provided list of `pmix_info_t` will be transferred to a `pmix_data_array_t` containing `pmix_info_t` structures.

Release a `pmix_info_t` list

Summary

Release the provided `pmix_info_t` list

1

Format

C

2

void

3

PMIx_Info_list_release(void *ptr);

C

4

IN ptr

5

A void* pointer initialized via `PMIx_Info_list_start` (handle)

6

Description

7

Information contained in the `pmix_info_t` on the list shall be released in addition to whatever backing storage the implementation may have allocated to support construction of the list.

8

9 3.2.10 Info Type Directives

10 *PMIx v2.0*

11

The `pmix_info_directives_t` structure is a `uint32_t` type that defines the behavior of command directives via `pmix_info_t` arrays. By default, the values in the `pmix_info_t` array passed to a PMIx are *optional*.

12

Advice to users

13

A PMIx implementation or PMIx-enabled RM may ignore any `pmix_info_t` value passed to a PMIx API that it does not support or does not recognize if it is not explicitly marked as `PMIX_INFO_REQD`. This is because the values specified default to optional, meaning they can be ignored in such circumstances. This may lead to unexpected behavior when porting between environments or PMIx implementations if the user is relying on the behavior specified by the `pmix_info_t` value. Users relying on the behavior defined by the `pmix_info_t` are advised to set the `PMIX_INFO_REQD` flag using the `PMIX_INFO_REQUIRED` macro.

14

15

16

17

18

19

Advice to PMIx library implementers

20

The top 16-bits of the `pmix_info_directives_t` are reserved for internal use by PMIx library implementers - the PMIx standard will *not* specify their intent, leaving them for customized use by implementers. Implementers are advised to use the provided `PMIX_INFO_IS_REQUIRED` macro for testing this flag, and must return `PMIX_ERR_NOT_SUPPORTED` as soon as possible to the caller if the required behavior is not supported.

21

22

23

24

1 The following constants were introduced in version 2.0 (unless otherwise marked) and can be used
2 to set a variable of the type `pmix_info_directives_t`.

3 **PMIX_INFO_REQD** `0x00000001` The behavior defined in the `pmix_info_t` array is
4 required, and not optional. This is a bit-mask value.

5 **PMIX_INFO_REQD_PROCESSED** `0x00000004` Mark that this required attribute has been
6 processed. A required attribute can be handled at any level - the PMIx client library might
7 take care of it, or it may be resolved by the PMIx server library, or it may pass up to the host
8 environment for handling. If a level does not recognize or support the required attribute, it is
9 required to pass it upwards to give the next level an opportunity to process it. Thus, the host
10 environment (or the server library if the host does not support the given operation) must know
11 if a lower level has handled the requirement so it can return a
12 **PMIX_ERR_NOT_SUPPORTED** error status if the host itself cannot meet the request. Upon
13 processing the request, the level must therefore mark the attribute with this directive to alert
14 any subsequent levels that the requirement has been met.

15 **PMIX_INFO_ARRAY_END** `0x00000002` Mark that this `pmix_info_t` struct is at the end
16 of an array created by the **PMIX_INFO_CREATE** macro. This is a bit-mask value.

17 **PMIX_INFO_DIR_RESERVED** `0xffff0000` A bit-mask identifying the bits reserved for
18 internal use by implementers - these currently are set as `0xffff0000`.

Advice to PMIx server hosts

19 Host environments are advised to use the provided **PMIX_INFO_IS_REQUIRED** macro for
20 testing this flag and must return **PMIX_ERR_NOT_SUPPORTED** as soon as possible to the caller if
21 the required behavior is not supported.

3.2.10.1 Info Directive support macros

22 The following macros are provided to support the setting and testing of `pmix_info_t` directives.

Mark an info structure as required

24 Set the **PMIX_INFO_REQD** flag in a `pmix_info_t` structure.

25 *PMIx v2.0*

```
26 PMIX_INFO_REQUIRED(info);
```

27 **IN** `info`

28 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

29 This macro simplifies the setting of the **PMIX_INFO_REQD** flag in `pmix_info_t` structures.

1 **Mark an info structure as optional**

2 Unsets the `PMIX_INFO_REQD` flag in a `pmix_info_t` structure.

▼ `C` ▼

3 `PMIX_INFO_OPTIONAL(info);`

▲ `C` ▲

4 **IN** `info`

5 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

6 This macro simplifies marking a `pmix_info_t` structure as *optional*.

7 **Test an info structure for *required* directive**

8 Test the `PMIX_INFO_REQD` flag in a `pmix_info_t` structure, returning `true` if the flag is set.

PMIx v2.0

▼ `C` ▼

9 `PMIX_INFO_IS_REQUIRED(info);`

▲ `C` ▲

10 **IN** `info`

11 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

12 This macro simplifies the testing of the required flag in `pmix_info_t` structures.

13 **Test an info structure for *optional* directive**

14 Test a `pmix_info_t` structure, returning `true` if the structure is *optional*.

PMIx v2.0

▼ `C` ▼

15 `PMIX_INFO_IS_OPTIONAL(info);`

▲ `C` ▲

16 **IN** `info`

17 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

18 Test the `PMIX_INFO_REQD` flag in a `pmix_info_t` structure, returning `true` if the flag is *not*
19 set.

20 **Mark a required attribute as processed**

21 Mark that a required `pmix_info_t` structure has been processed.

PMIx v4.0

▼ `C` ▼

22 `PMIX_INFO_PROCESSED(info);`

▲ `C` ▲

23 **IN** `info`

24 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

25 Set the `PMIX_INFO_REQD_PROCESSED` flag in a `pmix_info_t` structure indicating that is
26 has been processed.

1 **Test if a required attribute has been processed**

2 Test that a required `pmix_info_t` structure has been processed.

PMIx v4.0

```
▼ _____ C _____ ▼
```

3 `PMIX_INFO_WAS_PROCESSED(info);`

```
▲ _____ C _____ ▲
```

4 **IN** `info`

5 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

6 Test the `PMIX_INFO_REQD_PROCESSED` flag in a `pmix_info_t` structure.

7 **Test an info structure for *end of array* directive**

8 Test a `pmix_info_t` structure, returning `true` if the structure is at the end of an array created by
9 the `PMIX_INFO_CREATE` macro.

PMIx v2.2

```
▼ _____ C _____ ▼
```

10 `PMIX_INFO_IS_END(info);`

```
▲ _____ C _____ ▲
```

11 **IN** `info`

12 Pointer to the `pmix_info_t` (pointer to `pmix_info_t`)

13 This macro simplifies the testing of the end-of-array flag in `pmix_info_t` structures.

14 **3.2.11 Environmental Variable Structure**

PMIx v3.0

15 Define a structure for specifying environment variable modifications. Standard environment
16 variables (e.g., `PATH`, `LD_LIBRARY_PATH`, and `LD_PRELOAD`) take multiple arguments
17 separated by delimiters. Unfortunately, the delimiters depend upon the variable itself - some use
18 semi-colons, some colons, etc. Thus, the operation requires not only the name of the variable to be
19 modified and the value to be inserted, but also the separator to be used when composing the
20 aggregate value.

```
▼ _____ C _____ ▼
```

```
21 typedef struct {  
22     char *envar;  
23     char *value;  
24     char separator;  
25 } pmix_envar_t;
```

```
▲ _____ C _____ ▲
```

26 **3.2.11.1 Environmental variable support macros**

27 The following macros are provided to support the `pmix_envar_t` structure.

1 **Static initializer for the envar structure**

2 Provide a static initializer for the `pmix_envar_t` fields.



3 **PMIX_ENVAR_STATIC_INIT**



4 **Initialize the envar structure**

5 Initialize the `pmix_envar_t` fields.

PMIx v3.0



6 **PMIX_ENVAR_CONSTRUCT (m)**



7 **IN** m

8 Pointer to the structure to be initialized (pointer to `pmix_envar_t`)

9 **Destruct the envar structure**

10 Clear the `pmix_envar_t` fields.

PMIx v3.0



11 **PMIX_ENVAR_DESTRUCT (m)**



12 **IN** m

13 Pointer to the structure to be destructed (pointer to `pmix_envar_t`)

14 **Create an envar array**

15 Allocate and initialize an array of `pmix_envar_t` structures.

PMIx v3.0



16 **PMIX_ENVAR_CREATE (m, n)**



17 **INOUT** m

18 Address where the pointer to the array of `pmix_envar_t` structures shall be stored (handle)

19 **IN** n

20 Number of structures to be allocated (`size_t`)

1 **Free an envar array**
2 Release an array of `pmix_envar_t` structures.

▼ `PMIX_ENVAR_FREE` C ▼

3 **`PMIX_ENVAR_FREE`** (`m`, `n`)

▲ `PMIX_ENVAR_FREE` C ▲

4 **IN** `m`
5 Pointer to the array of `pmix_envar_t` structures (handle)
6 **IN** `n`
7 Number of structures in the array (`size_t`)

8 **Load an envar structure**
9 Load values into a `pmix_envar_t`.

PMIx v2.0 ▼ `PMIX_ENVAR_LOAD` C ▼

10 **`PMIX_ENVAR_LOAD`** (`m`, `e`, `v`, `s`)

▲ `PMIX_ENVAR_LOAD` C ▲

11 **IN** `m`
12 Pointer to the structure to be loaded (pointer to `pmix_envar_t`)
13 **IN** `e`
14 Environmental variable name (`char*`)
15 **IN** `v`
16 Value of variable (`char*`)
17 **IN** `v`
18 Separator character (`char`)

19 3.2.12 Byte Object Type

20 The `pmix_byte_object_t` structure describes a raw byte sequence.

PMIx v1.0 ▼ `pmix_byte_object_t` C ▼

```
21 typedef struct pmix_byte_object {  
22     char *bytes;  
23     size_t size;  
24 } pmix_byte_object_t;
```

▲ `pmix_byte_object_t` C ▲

25 3.2.12.1 Byte object support macros

26 The following macros support the `pmix_byte_object_t` structure.

1 **Static initializer for the byte object structure**

2 Provide a static initializer for the `pmix_byte_object_t` fields.



3 `PMIX_BYTE_OBJECT_STATIC_INIT`



4 **Initialize the byte object structure**

5 Initialize the `pmix_byte_object_t` fields.

PMIx v2.0



6 `PMIX_BYTE_OBJECT_CONSTRUCT (m)`



7 **IN** `m`

8 Pointer to the structure to be initialized (pointer to `pmix_byte_object_t`)

9 **Destruct the byte object structure**

10 Clear the `pmix_byte_object_t` fields.

PMIx v2.0



11 `PMIX_BYTE_OBJECT_DESTRUCT (m)`



12 **IN** `m`

13 Pointer to the structure to be destructed (pointer to `pmix_byte_object_t`)

14 **Create a byte object structure**

15 Allocate and initialize an array of `pmix_byte_object_t` structures.

PMIx v2.0



16 `PMIX_BYTE_OBJECT_CREATE (m, n)`



17 **INOUT** `m`

18 Address where the pointer to the array of `pmix_byte_object_t` structures shall be stored
19 (handle)

20 **IN** `n`

21 Number of structures to be allocated (`size_t`)

1 **Free a byte object array**
2 Release an array of `pmix_byte_object_t` structures.

▼ `C` ▼

3 **PMIX_BYTE_OBJECT_FREE**(m, n)

▲ `C` ▲

4 **IN** m
5 Pointer to the array of `pmix_byte_object_t` structures (handle)
6 **IN** n
7 Number of structures in the array (`size_t`)

8 **Load a byte object structure**
9 Load values into a `pmix_byte_object_t`.

PMIx v2.0 ▼ `C` ▼

10 **PMIX_BYTE_OBJECT_LOAD**(b, d, s)

▲ `C` ▲

11 **IN** b
12 Pointer to the structure to be loaded (pointer to `pmix_byte_object_t`)
13 **IN** d
14 Pointer to the data to be loaded (`char*`)
15 **IN** s
16 Number of bytes in the data array (`size_t`)

17 3.2.13 Data Array Structure

18 The `pmix_data_array_t` structure defines an array data structure.

PMIx v2.0 ▼ `C` ▼

```
19 typedef struct pmix_data_array {  
20     pmix_data_type_t type;  
21     size_t size;  
22     void *array;  
23 } pmix_data_array_t;
```

▲ `C` ▲

24 3.2.13.1 Data array support macros

25 The following macros support the `pmix_data_array_t` structure.

1 **Static initializer for the data array structure**

2 Provide a static initializer for the `pmix_data_array_t` fields.

▼ `PMIX_DATA_ARRAY_STATIC_INIT` C ▼

3 **PMIX_DATA_ARRAY_STATIC_INIT**

▲ `PMIX_DATA_ARRAY_STATIC_INIT` C ▲

4 **Initialize a data array structure**

5 Initialize the `pmix_data_array_t` fields, allocating memory for the array of the indicated type.

PMIx v2.2

▼ `PMIX_DATA_ARRAY_CONSTRUCT(m, n, t)` C ▼

6 **PMIX_DATA_ARRAY_CONSTRUCT(m, n, t)**

▲ `PMIX_DATA_ARRAY_CONSTRUCT(m, n, t)` C ▲

7 **IN m**

8 Pointer to the structure to be initialized (pointer to `pmix_data_array_t`)

9 **IN n**

10 Number of elements in the array (`size_t`)

11 **IN t**

12 PMIx data type of the array elements (`pmix_data_type_t`)

13 **Destruct a data array structure**

14 Destruct the `pmix_data_array_t`, releasing the memory in the array.

PMIx v2.2

▼ `PMIX_DATA_ARRAY_DESTRUCT(m)` C ▼

15 **PMIX_DATA_ARRAY_DESTRUCT(m)**

▲ `PMIX_DATA_ARRAY_DESTRUCT(m)` C ▲

16 **IN m**

17 Pointer to the structure to be destructed (pointer to `pmix_data_array_t`)

18 **Create a data array structure**

19 Allocate memory for the `pmix_data_array_t` object itself, and then allocate memory for the
20 array of the indicated type.

PMIx v2.2

▼ `PMIX_DATA_ARRAY_CREATE(m, n, t)` C ▼

21 **PMIX_DATA_ARRAY_CREATE(m, n, t)**

▲ `PMIX_DATA_ARRAY_CREATE(m, n, t)` C ▲

22 **INOUT m**

23 Variable to be set to the address of the structure (pointer to `pmix_data_array_t`)

24 **IN n**


25 Number of elements in the array (`size_t`)

26 **IN t**

27 PMIx data type of the array elements (`pmix_data_type_t`)

1 **Free a data array structure**

2 Release the memory in the array, and then release the `pmix_data_array_t` object itself.




6 **3.2.14 Argument Array Macros**

7 The following macros support the construction and release of **NULL**-terminated argv arrays of
8 strings.

9 **Argument array extension**

10 Append a string to a NULL-terminated, argv-style array of strings.



12 **OUT** `r`

13 Status code indicating success or failure of the operation (`pmix_status_t`)

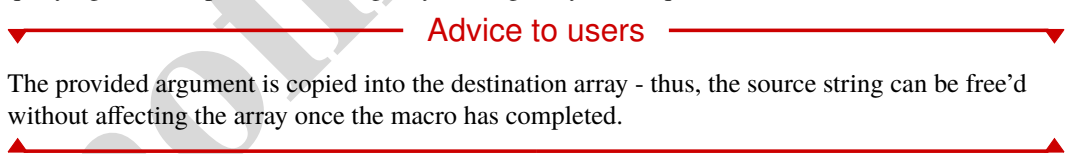
14 **INOUT** `a`

15 Argument list (pointer to NULL-terminated array of strings)

16 **IN** `b`

17 Argument to append to the list (string)

18 This function helps the caller build the **argv** portion of `pmix_app_t` structure, arrays of keys for
19 querying, or other places where argv-style string arrays are required.



Argument array prepend

Prepend a string to a NULL-terminated, argv-style array of strings.

▼ C ▼

```
PMIX_ARGV_PREPEND(r, a, b);
```

▲ C ▲

OUT r

Status code indicating success or failure of the operation ([pmix_status_t](#))

INOUT a

Argument list (pointer to NULL-terminated array of strings)

IN b

Argument to append to the list (string)

This function helps the caller build the **argv** portion of [pmix_app_t](#) structure, arrays of keys for querying, or other places where argv-style string arrays are required.

▼ Advice to users ▼

The provided argument is copied into the destination array - thus, the source string can be free'd without affecting the array once the macro has completed.

Argument array extension - unique

Append a string to a NULL-terminated, argv-style array of strings, but only if the provided argument doesn't already exist somewhere in the array.

▼ C ▼

```
PMIX_ARGV_APPEND_UNIQUE(r, a, b);
```

▲ C ▲

OUT r

Status code indicating success or failure of the operation ([pmix_status_t](#))

INOUT a

Argument list (pointer to NULL-terminated array of strings)

IN b

Argument to append to the list (string)

This function helps the caller build the **argv** portion of [pmix_app_t](#) structure, arrays of keys for querying, or other places where argv-style string arrays are required.

▼ Advice to users ▼

The provided argument is copied into the destination array - thus, the source string can be free'd without affecting the array once the macro has completed.

▲ ▲

1
2

Argument array release

Free an argv-style array and all of the strings that it contains.



3

```
PMIX_ARGV_FREE (a) ;
```



4
5

IN a

Argument list (pointer to NULL-terminated array of strings)

6

This function releases the array and all of the strings it contains.

7
8

Argument array split

Split a string into a NULL-terminated argv array.



9

```
PMIX_ARGV_SPLIT (a, b, c) ;
```



10

OUT a

Resulting argv-style array (**char****)

11

12

IN b

String to be split (**char***)

13

14

IN c

Delimiter character (**char**)

15

16

Split an input string into a NULL-terminated argv array. Do not include empty strings in the resulting array.

17



18

All strings are inserted into the argv array by value; the newly-allocated array makes no references to the src_string argument (i.e., it can be freed after calling this function without invalidating the output argv array)

19

20



1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26

Argument array join

Join all the elements of an argv array into a single newly-allocated string.

▼ C ▼

PMIX_ARGV_JOIN(a, b, c);

▲ C ▲

OUT a

Resulting string (**char***)

IN b

Argv-style array to be joined (**char****)

IN c

Delimiter character (**char**)

Join all the elements of an argv array into a single newly-allocated string.

Argument array count

Return the length of a NULL-terminated argv array.

▼ C ▼

PMIX_ARGV_COUNT(r, a);

▲ C ▲

OUT r

Number of strings in the array (integer)

IN a

Argv-style array (**char****)

Count the number of elements in an argv array

Argument array copy

Copy an argv array, including copying all of its strings.

▼ C ▼

PMIX_ARGV_COPY(a, b);

▲ C ▲

OUT a

New argv-style array (**char****)

IN b

Argv-style array (**char****)

Copy an argv array, including copying all of its strings.

1 3.2.15 Set Environment Variable

2 Summary

3 Set an environment variable in a **NULL**-terminated, env-style array.

```
4 PMIX_SETENV(r, name, value, env);
```

5 OUT *r*

6 Status code indicating success or failure of the operation ([pmix_status_t](#))

7 IN *name*

8 Argument name (string)

9 IN *value*

10 Argument value (string)

11 INOUT *env*

12 Environment array to update (pointer to array of strings)

13 Description

14 Similar to `setenv` from the C API, this allows the caller to set an environment variable in the
15 specified `env` array, which could then be passed to the [pmix_app_t](#) structure or any other
16 destination.

Advice to users

17 The provided name and value are copied into the destination environment array - thus, the source
18 strings can be free'd without affecting the array once the macro has completed.

19 3.3 Generalized Data Types Used for Packing/Unpacking

20 The [pmix_data_type_t](#) structure is a `uint16_t` type for identifying the data type for
21 packing/unpacking purposes. New data type values introduced in this version of the Standard are
22 shown in **magenta**.

Advice to PMIx library implementers

23 The following constants can be used to set a variable of the type [pmix_data_type_t](#). Data
24 types in the PMIx Standard are defined in terms of the C-programming language. Implementers
25 wishing to support other languages should provide the equivalent definitions in a
26 language-appropriate manner. Additionally, a PMIx implementation may choose to add additional
27 types.

1 **PMIX_UNDEF** 0 Undefined.
2 **PMIX_BOOL** 1 Boolean (converted to/from native **true/false**) (**bool**).
3 **PMIX_BYTE** 2 A byte of data (**uint8_t**).
4 **PMIX_STRING** 3 **NULL** terminated string (**char***).
5 **PMIX_SIZE** 4 Size **size_t**.
6 **PMIX_PID** 5 Operating Process Identifier (PID) (**pid_t**).
7 **PMIX_INT** 6 Integer (**int**).
8 **PMIX_INT8** 7 8-byte integer (**int8_t**).
9 **PMIX_INT16** 8 16-byte integer (**int16_t**).
10 **PMIX_INT32** 9 32-byte integer (**int32_t**).
11 **PMIX_INT64** 10 64-byte integer (**int64_t**).
12 **PMIX_UINT** 11 Unsigned integer (**unsigned int**).
13 **PMIX_UINT8** 12 Unsigned 8-byte integer (**uint8_t**).
14 **PMIX_UINT16** 13 Unsigned 16-byte integer (**uint16_t**).
15 **PMIX_UINT32** 14 Unsigned 32-byte integer (**uint32_t**).
16 **PMIX_UINT64** 15 Unsigned 64-byte integer (**uint64_t**).
17 **PMIX_FLOAT** 16 Float (**float**).
18 **PMIX_DOUBLE** 17 Double (**double**).
19 **PMIX_TIMEVAL** 18 Time value (**struct timeval**).
20 **PMIX_TIME** 19 Time (**time_t**).
21 **PMIX_STATUS** 20 Status code **pmix_status_t**.
22 **PMIX_VALUE** 21 Value (**pmix_value_t**).
23 **PMIX_PROC** 22 Process (**pmix_proc_t**).
24 **PMIX_APP** 23 Application context.
25 **PMIX_INFO** 24 Info object.
26 **PMIX_PDATA** 25 Pointer to data.
27 **PMIX_BYTE_OBJECT** 27 Byte object (**pmix_byte_object_t**).
28 **PMIX_KVAL** 28 Key/value pair.
29 **PMIX_PERSIST** 30 Persistence (**pmix_persistence_t**).
30 **PMIX_POINTER** 31 Pointer to an object (**void***).
31 **PMIX_SCOPE** 32 Scope (**pmix_scope_t**).
32 **PMIX_DATA_RANGE** 33 Range for data (**pmix_data_range_t**).
33 **PMIX_COMMAND** 34 PMIx command code (used internally).
34 **PMIX_INFO_DIRECTIVES** 35 Directives flag for **pmix_info_t**
35 (**pmix_info_directives_t**).
36 **PMIX_DATA_TYPE** 36 Data type code (**pmix_data_type_t**).
37 **PMIX_PROC_STATE** 37 Process state (**pmix_proc_state_t**).
38 **PMIX_PROC_INFO** 38 Process information (**pmix_proc_info_t**).
39 **PMIX_DATA_ARRAY** 39 Data array (**pmix_data_array_t**).
40 **PMIX_PROC_RANK** 40 Process rank (**pmix_rank_t**).
41 **PMIX_PROC_NAMESPACE** 60 Process namespace (**pmix_namespace_t**). %

1 **PMIX_QUERY** 41 Query structure ([pmix_query_t](#)).

2 **PMIX_COMPRESSED_STRING** 42 String compressed with zlib (**char***).

3 **Provisional** **PMIX_COMPRESSED_BYTE_OBJECT** 59 Byte object whose bytes have been compressed

4 with zlib ([pmix_byte_object_t](#)).

5 **PMIX_ALLOC_DIRECTIVE** 43 Allocation directive ([pmix_alloc_directive_t](#)).

6 **PMIX_IOF_CHANNEL** 45 Input/output forwarding channel ([pmix_iof_channel_t](#)).

7 **PMIX_ENVAR** 46 Environmental variable structure ([pmix_envar_t](#)).

8 **PMIX_COORD** 47 Structure containing fabric coordinates ([pmix_coord_t](#)).

9 **PMIX_REGATTR** 48 Structure supporting attribute registrations ([pmix_regattr_t](#)).

10 **PMIX_REGEX** 49 Regular expressions - can be a valid NULL-terminated string or an

11 arbitrary array of bytes.

12 **PMIX_JOB_STATE** 50 Job state ([pmix_job_state_t](#)).

13 **PMIX_LINK_STATE** 51 Link state ([pmix_link_state_t](#)).

14 **PMIX_PROC_CPUSSET** 52 Structure containing the binding bitmap of a process

15 ([pmix_cpuset_t](#)).

16 **PMIX_GEOMETRY** 53 Geometry structure containing the fabric coordinates of a specified

17 device. ([pmix_geometry_t](#)).

18 **PMIX_DEVICE_DIST** 54 Structure containing the minimum and maximum relative distance

19 from the caller to a given fabric device. ([pmix_device_distance_t](#)).

20 **PMIX_ENDPOINT** 55 Structure containing an assigned endpoint for a given fabric device.

21 ([pmix_endpoint_t](#)).

22 **PMIX_TOPO** 56 Structure containing the topology for a given node. ([pmix_topology_t](#)).

23 **PMIX_DEVTYPE** 57 Bitmask containing the types of devices being referenced.

24 ([pmix_device_type_t](#)).

25 **PMIX_LOCTYPE** 58 Bitmask describing the relative location of another process.

26 ([pmix_locality_t](#)).

27 **PMIX_STOR_MEDIUM** 66 Bitmask specifying different types of storage mediums.

28 ([pmix_storage_medium_t](#)).

29 **PMIX_STOR_ACCESS** 67 Bitmask specifying different levels of storage accessibility (i.e.,

30 from where a storage system may be accessed). ([pmix_storage_accessibility_t](#)).

31 **PMIX_STOR_PERSIST** 68 Bitmask specifying different levels of persistence for a particular

32 storage system. ([pmix_storage_persistence_t](#)).

33 **PMIX_STOR_ACCESS_TYPE** 69 Bitmask specifying different storage system access types.

34 ([pmix_storage_access_type_t](#)).

35 **PMIX_DATA_TYPE_MAX** A starting point for implementer-specific data types. Values above

36 this are guaranteed not to conflict with PMIX values. Definitions should always be based on

37 the **PMIX_DATA_TYPE_MAX** constant and not a specific value as the value of the constant

38 may change.

1 3.4 General Callback Functions

2 PMIx provides blocking and nonblocking versions of most APIs. In the nonblocking versions, a
3 callback is activated upon completion of the the operation. This section describes many of those
4 callbacks.

5 3.4.1 Release Callback Function

6 Summary

7 The `pmix_release_cbfunc_t` is used by the `pmix_modex_cbfunc_t` and
8 `pmix_info_cbfunc_t` operations to indicate that the callback data may be reclaimed/freed by
9 the caller.

10 *PMIx v1.0*

Format

```
typedef void (*pmix_release_cbfunc_t)
    (void *cbdata);
```

13 INOUT cbdata

14 Callback data passed to original API call (memory reference)

15 Description

16 Since the data is “owned” by the host server, provide a callback function to notify the host server
17 that we are done with the data so it can be released.

18 3.4.2 Lookup Callback Function

19 Summary

20 The `pmix_lookup_cbfunc_t` is used by `PMIx_Lookup_nb` to return data.

21 *PMIx v1.0*

```
typedef void (*pmix_lookup_cbfunc_t)
    (pmix_status_t status,
     pmix_pdata_t data[], size_t ndata,
     void *cbdata);
```

25 IN status

26 Status associated with the operation (handle)

27 IN data

28 Array of data returned (`pmix_pdata_t`)

29 IN ndata

30 Number of elements in the `data` array (`size_t`)

31 IN cbdata

32 Callback data passed to original API call (memory reference)

Description

A callback function for calls to `PMIx_Lookup_nb`. The function will be called upon completion of the `PMIx_Lookup_nb` API with the `status` indicating the success or failure of the request. Any retrieved data will be returned in an array of `pmix_pdata_t` structs. The namespace and rank of the process that provided each data element is also returned.

Note that the `pmix_pdata_t` structures will be released upon return from the callback function, so the receiver must copy/protect the data prior to returning if it needs to be retained.

3.4.3 Op Callback Function

Summary

The `pmix_op_cbfunc_t` is used by operations that simply return a status.

PMIx v1.0

```
typedef void (*pmix_op_cbfunc_t)
    (pmix_status_t status, void *cbdata);
```

IN status

Status associated with the operation (handle)

IN cbdata

Callback data passed to original API call (memory reference)

Description

Used by a wide range of PMIx API's including `PMIx_Fence_nb`, `pmix_server_client_connected2_fn_t`, `PMIx_server_register_nspace`. This callback function is used to return a status to an often nonblocking operation.

3.4.4 Value Callback Function

Summary

The `pmix_value_cbfunc_t` is used by `PMIx_Get_nb` to return data.

PMIx v1.0

```
typedef void (*pmix_value_cbfunc_t)
    (pmix_status_t status,
     pmix_value_t *kv, void *cbdata);
```

IN status

Status associated with the operation (handle)

IN kv

Key/value pair representing the data (`pmix_value_t`)

IN cbdata

Callback data passed to original API call (memory reference)

Description

A callback function for calls to `PMIx_Get_nb`. The *status* indicates if the requested data was found or not. A pointer to the `pmix_value_t` structure containing the found data is returned. The pointer will be `NULL` if the requested data was not found.

3.4.5 Info Callback Function

Summary

The `pmix_info_cbfunc_t` is a general information callback used by various APIs.

PMIx v2.0

```
typedef void (*pmix_info_cbfunc_t)
    (pmix_status_t status,
     pmix_info_t info[], size_t ninfo,
     void *cbdata,
     pmix_release_cbfunc_t release_fn,
     void *release_cbdata);
```

IN status

Status associated with the operation (`pmix_status_t`)

IN info

Array of `pmix_info_t` returned by the operation (pointer)

IN ninfo

Number of elements in the *info* array (`size_t`)

IN cbdata

Callback data passed to original API call (memory reference)

IN release_fn

Function to be called when done with the *info* data (function pointer)

IN release_cbdata

Callback data to be passed to *release_fn* (memory reference)

Description

The *status* indicates if requested data was found or not. An array of `pmix_info_t` will contain the key/value pairs.

3.4.6 Handler registration callback function

Summary

Callback function for calls to register handlers, e.g., event notification and IOF requests.

1 **Format** C

```
2 typedef void (*pmix_hdlr_reg_cbfunc_t)
3     (pmix_status_t status,
4      size_t refid,
5      void *cbdata);
```

C

- 6 **IN** **status**
7 **PMIX_SUCCESS** or an appropriate error constant (**pmix_status_t**)
- 8 **IN** **refid**
9 reference identifier assigned to the handler by PMIx, used to deregister the handler (**size_t**)
- 10 **IN** **cbdata**
11 object provided to the registration call (pointer)

12 **Description**
13 Callback function for calls to register handlers, e.g., event notification and IOF requests.

14 3.5 PMIx Datatype Value String Representations

15 Provide a string representation for several types of values. Note that the provided string is statically
16 defined and must NOT be **free**'d.

17 **Summary**
18 String representation of a **pmix_status_t**.

19 *PMIx v1.0* C

```
19 const char*
20 PMIx_Error_string(pmix_status_t status);
```

21 **Summary**
22 String representation of a **pmix_proc_state_t**.

23 *PMIx v2.0* C

```
23 const char*
24 PMIx_Proc_state_string(pmix_proc_state_t state);
```


1 **Summary**
2 String representation of a [pmix_scope_t](#).

▼ [C](#) [_____](#) ▼

3 **const char***
4 **PMIx_Scope_string**(pmix_scope_t scope);

▲ [C](#) [_____](#) ▲

5 **Summary**
6 String representation of a [pmix_persistence_t](#).

PMIx v2.0 ▼ [C](#) [_____](#) ▼

7 **const char***
8 **PMIx_Persistence_string**(pmix_persistence_t persist);

▲ [C](#) [_____](#) ▲

9 **Summary**
10 String representation of a [pmix_data_range_t](#).

PMIx v2.0 ▼ [C](#) [_____](#) ▼

11 **const char***
12 **PMIx_Data_range_string**(pmix_data_range_t range);

▲ [C](#) [_____](#) ▲

13 **Summary**
14 String representation of a [pmix_info_directives_t](#).

PMIx v2.0 ▼ [C](#) [_____](#) ▼

15 **const char***
16 **PMIx_Info_directives_string**(pmix_info_directives_t directives);

▲ [C](#) [_____](#) ▲

17 **Summary**
18 String representation of a [pmix_data_type_t](#).

PMIx v2.0 ▼ [C](#) [_____](#) ▼

19 **const char***
20 **PMIx_Data_type_string**(pmix_data_type_t type);

▲ [C](#) [_____](#) ▲

1 **Summary**
2 String representation of a `pmix_alloc_directive_t`.

▼ C ▼

3 `const char*`
4 `PMIx_Alloc_directive_string(pmix_alloc_directive_t directive);`

▲ C ▲

5 **Summary**
6 String representation of a `pmix_iof_channel_t`.

PMIx v3.0 ▼ C ▼

7 `const char*`
8 `PMIx_IOF_channel_string(pmix_iof_channel_t channel);`

▲ C ▲

9 **Summary**
10 String representation of a `pmix_job_state_t`.

PMIx v4.0 ▼ C ▼

11 `const char*`
12 `PMIx_Job_state_string(pmix_job_state_t state);`

▲ C ▲

13 **Summary**
14 String representation of a PMIx attribute.

PMIx v4.0 ▼ C ▼

15 `const char*`
16 `PMIx_Get_attribute_string(char *attributename);`

▲ C ▲

17 **Summary**
18 Return the PMIx attribute name corresponding to the given attribute string.

PMIx v4.0 ▼ C ▼

19 `const char*`
20 `PMIx_Get_attribute_name(char *attributestring);`

▲ C ▲

1 **Summary**
2 String representation of a `pmix_link_state_t`.



3 `const char*`
4 `PMIx_Link_state_string(pmix_link_state_t state);`



5 **Summary**
6 String representation of a `pmix_device_type_t`.

PMIx v4.0



7 `const char*`
8 `PMIx_Device_type_string(pmix_device_type_t type);`



Unofficial Draft

CHAPTER 4

Client Initialization and Finalization

1 The PMIx library is required to be initialized and finalized around the usage of most PMIx
2 functions or macros. The APIs that may be used outside of the initialized and finalized region are
3 noted. All other APIs must be used inside this region.

4 There are three sets of initialization and finalization functions depending upon the role of the
5 process in the PMIx Standard - those associated with the PMIx *client* are defined in this chapter.
6 Similar functions corresponding to the roles of *server* and *tool* are defined in Chapters 17 and 18,
7 respectively.

8 Note that a process can only call *one* of the initialization/finalization functional pairs from the set of
9 three - e.g., a process that calls the client initialization function cannot also call the tool or server
10 initialization functions, and must call the corresponding client finalization function. Regardless of
11 the role assumed by the process, all processes have access to the client APIs. Thus, the *server* and
12 *tool* roles can be considered supersets of the PMIx client.

4.1 PMIx_Initialized

Summary

14 Determine if the PMIx library has been initialized. This function may be used outside of the
15 initialized and finalized region, and is usable by servers and tools in addition to clients.
16

Format

17 *PMIx v1.0*

18 `int PMIx_Initialized(void)`

19 A value of **1** (true) will be returned if the PMIx library has been initialized, and **0** (false) otherwise.

Rationale

20 The return value is an integer for historical reasons as that was the signature of prior PMI libraries.

Description

21 Check to see if the PMIx library has been initialized using any of the init functions: `PMIx_Init`,
22 `PMIx_server_init`, or `PMIx_tool_init`.
23

1 4.2 PMIx_Get_version

2 Summary

3 Get the PMIx version information. This function may be used outside of the initialized and
4 finalized region, and is usable by servers and tools in addition to clients.

5 Format

PMIx v1.0

C

6 `const char* PMIx_Get_version(void)`

C

7 Description

8 Get the PMIx version string. Note that the provided string is statically defined and must *not* be
9 free'd.

10 4.3 PMIx_Init

11 Summary

12 Initialize the PMIx client library

13 Format

PMIx v1.2

C

14 `pmix_status_t`

15 `PMIx_Init(pmix_proc_t *proc,`
16 `pmix_info_t info[], size_t ninfo)`

C

17 INOUT `proc`

18 `proc` structure (handle)

19 IN `info`

20 Array of `pmix_info_t` structures (array of handles)

21 IN `ninfo`

22 Number of elements in the `info` array (`size_t`)

23 Returns `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant.

Optional Attributes

The following attributes are optional for implementers of PMIx libraries:

PMIX_USOCK_DISABLE "pmix.usock.disable" (bool)

Disable legacy UNIX socket (usock) support. If the library supports Unix socket connections, this attribute may be supported for disabling it.

PMIX_SOCKET_MODE "pmix.sockmode" (uint32_t)

POSIX *mode_t* (9 bits valid). If the library supports socket connections, this attribute may be supported for setting the socket mode.

PMIX_SINGLE_LISTENER "pmix.sing.listnr" (bool)

Use only one rendezvous socket, letting priorities and/or environment parameters select the active transport. If the library supports multiple methods for clients to connect to servers, this attribute may be supported for disabling all but one of them.

PMIX_TCP_REPORT_URI "pmix.tcp.repuri" (char*)

If provided, directs that the TCP Uniform Resource Identifier (URI) be reported and indicates the desired method of reporting: '-' for stdout, '+' for stderr, or filename. If the library supports TCP socket connections, this attribute may be supported for reporting the URI.

PMIX_TCP_IF_INCLUDE "pmix.tcp.ifinclude" (char*)

Comma-delimited list of devices and/or Classless Inter-Domain Routing (CIDR) notation to include when establishing the TCP connection. If the library supports TCP socket connections, this attribute may be supported for specifying the interfaces to be used.

PMIX_TCP_IF_EXCLUDE "pmix.tcp.ifexclude" (char*)

Comma-delimited list of devices and/or CIDR notation to exclude when establishing the TCP connection. If the library supports TCP socket connections, this attribute may be supported for specifying the interfaces that are *not* to be used.

PMIX_TCP_IPV4_PORT "pmix.tcp.ipv4" (int)

The IPv4 port to be used. If the library supports IPV4 connections, this attribute may be supported for specifying the port to be used.

PMIX_TCP_IPV6_PORT "pmix.tcp.ipv6" (int)

The IPv6 port to be used. If the library supports IPV6 connections, this attribute may be supported for specifying the port to be used.

PMIX_TCP_DISABLE_IPV4 "pmix.tcp.disipv4" (bool)

Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections, this attribute may be supported for disabling it.

PMIX_TCP_DISABLE_IPV6 "pmix.tcp.disipv6" (bool)

Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections, this attribute may be supported for disabling it.

PMIX_EXTERNAL_PROGRESS "pmix.evext" (bool)

1 The host shall progress the PMIx library via calls to **PMIx_Progress**

2 **PMIX_EVENT_BASE** "pmix.evbase" (void*)

3 Pointer to an **event_base** to use in place of the internal progress thread. All PMIx library
4 events are to be assigned to the provided event base. The event base *must* be compatible with
5 the event library used by the PMIx implementation - e.g., either both the host and PMIx
6 library must use libevent, or both must use libev. Cross-matches are unlikely to work and
7 should be avoided - it is the responsibility of the host to ensure that the PMIx
8 implementation supports (and was built with) the appropriate event library.

9 If provided, the following attributes are used by the event notification system for inter-library
10 coordination:

11 **PMIX_PROGRAMMING_MODEL** "pmix.pgm.model" (char*)

12 Programming model being initialized (e.g., "MPI" or "OpenMP").

13 **PMIX_MODEL_LIBRARY_NAME** "pmix.mdl.name" (char*)

14 Programming model implementation ID (e.g., "OpenMPI" or "MPICH").

15 **PMIX_MODEL_LIBRARY_VERSION** "pmix.mld.vrs" (char*)

16 Programming model version string (e.g., "2.1.1").

17 **PMIX_THREADING_MODEL** "pmix.threads" (char*)

18 Threading model used (e.g., "pthreads").

19 **PMIX_MODEL_NUM_THREADS** "pmix.mdl.nthrds" (uint64_t)

20 Number of active threads being used by the model.

21 **PMIX_MODEL_NUM_CPUS** "pmix.mdl.ncpu" (uint64_t)

22 Number of cpus being used by the model.

23 **PMIX_MODEL_CPU_TYPE** "pmix.mdl.cputype" (char*)

24 Granularity - "hwthread", "core", etc.

25 **PMIX_MODEL_AFFINITY_POLICY** "pmix.mdl.tap" (char*)

26 Thread affinity policy - e.g.: "master" (thread co-located with master thread), "close" (thread
27 located on cpu close to master thread), "spread" (threads load-balanced across available
28 cpus).



Description

Initialize the PMIx client, returning the process identifier assigned to this client's application in the provided `pmix_proc_t` struct. Passing a value of `NULL` for this parameter is allowed if the user wishes solely to initialize the PMIx system and does not require return of the identifier at that time.

When called, the PMIx client shall check for the required connection information of the local PMIx server and establish the connection. If the information is not found, or the server connection fails, then an appropriate error constant shall be returned.

If successful, the function shall return `PMIX_SUCCESS` and fill the `proc` structure (if provided) with the server-assigned namespace and rank of the process within the application. In addition, all startup information provided by the resource manager shall be made available to the client process via subsequent calls to `PMIx_Get`.

The PMIx client library shall be reference counted, and so multiple calls to `PMIx_Init` are allowed by the standard. Thus, one way for an application process to obtain its namespace and rank is to simply call `PMIx_Init` with a non-`NULL` `proc` parameter. Note that each call to `PMIx_Init` must be balanced with a call to `PMIx_Finalize` to maintain the reference count.

Each call to `PMIx_Init` may contain an array of `pmix_info_t` structures passing directives to the PMIx client library as per the above attributes.

Multiple calls to `PMIx_Init` shall not include conflicting directives. The `PMIx_Init` function will return an error when directives that conflict with prior directives are encountered.

4.3.1 Initialization events

The following events are typically associated with calls to `PMIx_Init`:

<code>PMIX_MODEL_DECLARED</code>	<code>-147</code>	Model declared.
<code>PMIX_MODEL_RESOURCES</code>	<code>-151</code>	Resource usage by a programming model has changed.
<code>PMIX_OPENMP_PARALLEL_ENTERED</code>	<code>-152</code>	An OpenMP parallel code region has been entered.
<code>PMIX_OPENMP_PARALLEL_EXITED</code>	<code>-153</code>	An OpenMP parallel code region has completed.

4.3.2 Initialization attributes

The following attributes influence the behavior of `PMIx_Init`.

4.3.2.1 Connection attributes

These attributes are used to describe a TCP socket for rendezvous with the local RM by passing them into the relevant initialization API - thus, they are not typically accessed via the `PMIx_Get` API.

`PMIX_TCP_REPORT_URI` "pmix.tcp.repuri" (`char*`)

1 If provided, directs that the TCP URI be reported and indicates the desired method of
2 reporting: '-' for stdout, '+' for stderr, or filename.

3 **PMIX_TCP_URI** "pmix.tcp.uri" (char*)

4 The URI of the PMIx server to connect to, or a file name containing it in the form of
5 file:<name of file containing it>.

6 **PMIX_TCP_IF_INCLUDE** "pmix.tcp.ifinclude" (char*)

7 Comma-delimited list of devices and/or CIDR notation to include when establishing the
8 TCP connection.

9 **PMIX_TCP_IF_EXCLUDE** "pmix.tcp.ifexclude" (char*)

10 Comma-delimited list of devices and/or CIDR notation to exclude when establishing the
11 TCP connection.

12 **PMIX_TCP_IPV4_PORT** "pmix.tcp.ipv4" (int)

13 The IPv4 port to be used..

14 **PMIX_TCP_IPV6_PORT** "pmix.tcp.ipv6" (int)

15 The IPv6 port to be used.

16 **PMIX_TCP_DISABLE_IPV4** "pmix.tcp.disipv4" (bool)

17 Set to **true** to disable IPv4 family of addresses.

18 **PMIX_TCP_DISABLE_IPV6** "pmix.tcp.disipv6" (bool)

19 Set to **true** to disable IPv6 family of addresses.

20 4.3.2.2 Programming model attributes

21 These attributes are associated with programming models.

22 **PMIX_PROGRAMMING_MODEL** "pmix.pgm.model" (char*)

23 Programming model being initialized (e.g., "MPI" or "OpenMP").

24 **PMIX_MODEL_LIBRARY_NAME** "pmix.mdl.name" (char*)

25 Programming model implementation ID (e.g., "OpenMPI" or "MPICH").

26 **PMIX_MODEL_LIBRARY_VERSION** "pmix.mld.vrs" (char*)

27 Programming model version string (e.g., "2.1.1").

28 **PMIX_THREADING_MODEL** "pmix.threads" (char*)

29 Threading model used (e.g., "pthreads").

30 **PMIX_MODEL_NUM_THREADS** "pmix.mdl.nthrds" (uint64_t)

31 Number of active threads being used by the model.

32 **PMIX_MODEL_NUM_CPUS** "pmix.mdl.ncpu" (uint64_t)

33 Number of cpus being used by the model.

34 **PMIX_MODEL_CPU_TYPE** "pmix.mdl.cputype" (char*)

35 Granularity - "hwthread", "core", etc.

36 **PMIX_MODEL_PHASE_NAME** "pmix.mdl.phase" (char*)

37 User-assigned name for a phase in the application execution (e.g., "cfd reduction").

38 **PMIX_MODEL_PHASE_TYPE** "pmix.mdl.ptype" (char*)

39 Type of phase being executed (e.g., "matrix multiply").

40 **PMIX_MODEL_AFFINITY_POLICY** "pmix.mdl.tap" (char*)

41 Thread affinity policy - e.g.: "master" (thread co-located with master thread), "close" (thread
42 located on cpu close to master thread), "spread" (threads load-balanced across available
43 cpus).

1 4.4 PMIx_Finalize

2 Summary

3 Finalize the PMIx client library.

4 *PMIx v1.0* Format

C

5 `pmix_status_t`

6 `PMIx_Finalize(const pmix_info_t info[], size_t ninfo)`

C

7 **IN** `info`

8 Array of `pmix_info_t` structures (array of handles)

9 **IN** `ninfo`

10 Number of elements in the `info` array (`size_t`)

11 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Optional Attributes

12 The following attributes are optional for implementers of PMIx libraries:

13 `PMIX_EMBED_BARRIER` "`pmix.embed.barrier`" (`bool`)

14 Execute a blocking fence operation before executing the specified operation.

15 `PMIx_Finalize` does not include an internal barrier operation by default. This attribute
16 directs `PMIx_Finalize` to execute a barrier as part of the finalize operation.

17 Description

18 Decrement the PMIx client library reference count. When the reference count reaches zero, the
19 library will finalize the PMIx client, closing the connection with the local PMIx server and
20 releasing all internally allocated memory.

21 4.4.1 Finalize attributes

22 The following attribute influences the behavior of `PMIx_Finalize`.

23 `PMIX_EMBED_BARRIER` "`pmix.embed.barrier`" (`bool`)

24 Execute a blocking fence operation before executing the specified operation.

25 `PMIx_Finalize` does not include an internal barrier operation by default. This attribute
26 directs `PMIx_Finalize` to execute a barrier as part of the finalize operation.

27 4.5 PMIx_Progress

28 Summary

29 Progress the PMIx library.

1
2
3
4
5
6

Format

C

void
PMIx_Progress (void)

C

Description

Progress the PMIx library. Note that special care must be taken to avoid deadlocking in PMIx callback functions and APIs.

Unofficial Draft

CHAPTER 5

Data Access and Sharing

1 In PMIx key-value pairs are the primary way that information is shared between processes in the
2 PMIx universe. A key-value pair consists of a string key and a value data structure which includes
3 the type of the value's data. Each key is unique within the context or domain in which it exists. The
4 context or domain of a key in PMIx is called a data realm. This chapter describes how key-value
5 pairs are made available and accessed by clients.

6 A key-value pair has a scope defined by its data realm. PMIx defines data realms for sessions, jobs,
7 applications, nodes, processes and a special realm for publish/lookup information. These are
8 explained more fully in chapters 6 and 9.

9 The source or provider of key-value pairs can be the host environment or a PMIx client. The host
10 environment can provide key-values pairs in any realm except for the publish/lookup realm. PMIx
11 clients can only provide key-values pairs within their own process realm and the publish/lookup
12 realm. The value of a key is not necessarily static and can be modified by the original provider of
13 the key.

14 PMIx classifies keys as either reserved or non-reserved. Only keys provided by a host environment
15 or PMIx implementation can begin with "**pmix**" and are called reserved keys. Host environments
16 and PMIx implementations must only provide keys that begin with "**pmix**" to avoid conflicts with
17 keys that are provided by clients. Reserved keys are documented in Chapter 6. Each standardized
18 reserved key has an associated macro that is documented along with its description. Likewise
19 clients can only provide keys that do not begin with "**pmix**" to avoid conflicting with keys
20 provided by the host environment or PMIx implementation. The string representation and the
21 macro are equivalent methods to refer to a key. Throughout this document, the macro is used rather
22 than its string value to refer to reserved keys. Clients cannot create reserved keys, nor can they
23 modify the value of reserved keys. The way reserved keys are accessed by clients is different than
24 for non-reserved keys as explained in Chapter 6.

25 Host environments are not required to provide all keys that are documented in the standard.
26 However, when providing keys that are documented, they must adhere to the behavior defined by
27 the standard for reserved keys. An implementation cannot provide a reserved key that behaves
28 differently than what is in the standard. Implementations are encouraged to standardize any
29 non-standardized, reserved keys that they provide to avoid conflicting with other implementations
30 or efforts to standardize the same key.

31 Host environments are permitted to provide non-reserved keys and non-standardized reserved keys.
32 However there are important caveats to providing either of these categories of keys that are
33 explained in Section 5.1 and Chapter 6. Host environments should avoid providing reserved keys

1 which are not standardized. If the reserved key is later standardized to behave in a way that is
2 different from the behavior provided by the host environment, the host environment will not be
3 compliant with the PMIx standard.

4 Reserved keys are set by the host environment and are made available at client initialization. Each
5 key is associated with a data realm and the retrieval API call will indicate the realm being queried.
6 Reserved keys are used to access information about a client's execution context. After initialization,
7 the client can access reserved keys using the **PMIx_Get** API. Common information that a client
8 may access includes, but is not limited to, the name of the host it is executing on, the number of
9 peer processes in its namespace, the number of peer processes in its application, and its job rank.

10 Non-reserved keys are provided by clients. A call to **PMIx_Put** will result in a key-value being
11 associated with the calling client within its process data realm. There is no mechanism to retract a
12 key-value that has been made available by a client through **PMIx_Put**. The key can be assigned a
13 new value, but cannot be removed entirely. Although key values can originate from different
14 sources, they are always retrieved using the **PMIx_Get** and **PMIx_Get_nb** API.

15 The publish/lookup data realm is accessed through a separate set of APIs. **PMIx_Put** cannot add
16 or modify key-values within the publish/lookup realm and **PMIx_Lookup** cannot access
17 key-values outside the publish/lookup realm. This data realm is described in detail in chapter 9.
18 Although **PMIx_Publish** and **PMIx_Lookup** are analogous to **PMIx_Put** and **PMIx_Get** in
19 that both pairs of APIs provide functionality for exposing and retrieving key-values, the semantics
20 vary significantly. For example, **PMIx_Lookup** includes the ability to wait until a key is defined
21 before returning. **PMIx_Publish** can restrict the scope of who can access data to its own
22 namespace, the host environment, the session it is executing in, all processes or even custom
23 scopes, while **PMIx_Put** can restrict the scope of who can access data to the node it is executing
24 on, remote nodes, all nodes or only itself. The publish/lookup data realm is useful for advertising
25 information that is not necessarily specific to one process to other processes in the PMIx universe.
26 The process accessing this information does *not* need to know the identity of the process that
27 provided the data.

28 PMIx does not provide a mechanism to asynchronously notify a process about the availability of
29 key-value information once it is made available by another process. However, the nonblocking
30 accessor interfaces (e.g., **PMIx_Get_nb**, **PMIx_Lookup_nb**) may provide a degree of
31 asynchronous notification on information availability.

32 Process related key-value exchanges allow a PMIx process to share information specific to itself,
33 and access information specific to one or more processes in the PMIx universe. These interactions
34 occur within the process data realm. The 'put/commit/get' exchange pattern is often used to
35 exchange process related information. Optionally, a 'put/commit/fence/get' exchange pattern adds
36 the 'fence' synchronization (and possible collective exchange) for applications that desire it.
37 Commonly, these exchange patterns are used in a *business card exchange* (a.k.a. *modex exchange*)
38 where one PMIx client shares its connectivity information, then other PMIx clients access that
39 information to establish a connection with that client. In some environments that support
40 "instant-on" all connectivity information for PMIx clients is stored in the job-level information at
41 process creation time and is accessible to the clients without the need to perform any additional

1 key-value exchange.

2 Keys can exist in multiple data realms, possibly with different values. PMIx clients can access
3 available information associated with any of the data realms. For example, a client can access the
4 number of nodes (**PMIX_NUM_NODES**) used by a session, job, or application. Rather than having
5 three different attributes, a single attribute is used but with the data realm context of the query
6 specified as additional attributes. Examples of these access patterns are presented in Section 6.2.1.

7 PMIx clients can share key-value pairs associated with themselves by using the **PMIx_Put**
8 function. The **PMIx_Put** function automatically associates the key-value pair with the calling
9 process, thus making it specific to that process. A client may call **PMIx_Put** as many times as
10 necessary and the data is not available to other processes until explicitly committed. A client must
11 call **PMIx_Commit** to make accessible all key-value pairs previously put by this process to all
12 other processes in the PMIx universe. This put and commit pattern provides implementors the
13 opportunity to make individual **PMIx_Put** calls efficient local operations, and then make the
14 whole set of key-value pairs accessible in a single step.

15 PMIx clients can access the key-value pairs associated with any process data realm in the PMIx
16 universe (including the calling process) by passing the specific process name of the target process
17 to the **PMIx_Get** and **PMIx_Get_nb** functions. The PMIx server local to the calling process will
18 retrieve that key-value pair from the PMIx server associated with the target process. Clients can
19 also access session, job, application, node, and namespace level information by using the
20 **PMIx_Get** and **PMIx_Get_nb** functions as shown in Section 6.2.1. The completion semantics
21 for **PMIx_Get** and **PMIx_Get_nb** differ depending on the type of key and its availability. See
22 Sections 5.3 and 6.2. For example, if a non-reserved key is not available, the **PMIx_Get** or
23 **PMIx_Get_nb** call will not complete, by default, until that key-value pair becomes available.

24 5.1 Non-reserved keys

25 *Non-reserved keys* are keys whose string representation begin with a prefix other than "**pmix**".
26 Such keys are defined by a client when information needs to be exchanged between processes, for
27 example, where connection information is required and the host environment does not support the
28 *instant on* option or where the host environment does not provide a required piece of data. Other
29 than the prefix, there are no restrictions on the use or content of non-reserved keys.

30 PMIx provides support for two methods of exchanging non-reserved keys:

- 31 • Global, collective exchange of the information prior to retrieval. This is accomplished by
32 executing a barrier operation that includes collection and exchange of the data provided by each
33 process such that each process has access to the full set of data from all participants once the
34 operation has completed. PMIx provides the **PMIx_Fence** function (or its non-blocking
35 equivalent) for this purpose, accompanied by the **PMIX_COLLECT_DATA** qualifier.
- 36 • Direct, on-demand retrieval of the information. No barrier or global exchange is conducted in
37 this case. Instead, information is retrieved from the host where that process is executing upon
38 request - i.e., a call to **PMIx_Get** results in a data exchange with the PMIx server on the remote

1 host. Various caching strategies may be employed by the host environment and/or PMIx
2 implementation to reduce the number of retrievals. Note that this method requires that the host
3 environment both know the location of the posting process and support direct information
4 retrieval.

5 Both of the above methods are based on retrieval from a specific process - i.e., the *proc* argument to
6 [PMIx_Get](#) must include both the namespace and the rank of the process that posted the
7 information. However, in some cases, non-reserved keys are provided on a globally unique basis
8 and the retrieving process has no knowledge of the identity of the process posting the key. This is
9 typically found in legacy applications (where the originating process identifier is often embedded in
10 the key itself) and in unstructured applications that lack rank-related behavior. In these cases, the
11 key remains associated with the namespace of the process that posted it, but is retrieved by use of
12 the [PMIX_RANK_UNDEF](#) rank. In addition, the keys must be globally exchanged prior to retrieval
13 as there is no way for the host to otherwise locate the source for the information.

14 Note that the retrieval rules for non-reserved keys (detailed in Section 5.3) differ significantly from
15 those used for reserved keys.

16 5.2 Posting Key/Value Pairs

17 PMIx clients can post non-reserved key-value pairs associated with themselves by using
18 [PMIx_Put](#). Alternatively, PMIx clients can cache both reserved and non-reserved key-value pairs
19 accessible only by the caller via the [PMIx_Store_internal](#) API.

20 5.2.1 PMIx_Put

21 Summary

22 Stage a key/value pair in preparation for being made accessible to processes.

23 Format

PMIx v1.0

C

```
24 pmix_status_t  
25 PMIx_Put(pmix_scope_t scope,  
26         const pmix_key_t key,  
27         pmix_value_t *val);
```

C

28 **IN** **scope**
29 Distribution scope of the provided value (handle)
30 **IN** **key**
31 key ([pmix_key_t](#))
32 **IN** **value**
33 Reference to a [pmix_value_t](#) structure (handle)

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- `PMIX_ERR_BAD_PARAM` indicating a reserved key is provided in the *key* argument.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Description

Post a key-value pair for distribution. Depending upon the PMIx implementation, the posted value may be locally cached in the client's PMIx library until `PMIx_Commit` is called.

The provided *scope* determines the ability of other processes to access the posted data, as defined in Section 5.2.1.1 on page 75. Specific implementations may support different scope values, but all implementations must support at least `PMIX_GLOBAL`.

The `pmix_value_t` structure supports both string and binary values. PMIx implementations are required to support heterogeneous environments by properly converting binary values between host architectures, and will copy the provided *value* into internal memory prior to returning from `PMIx_Put`.

Advice to users

Note that keys starting with a string of “`pmix`” must not be used in calls to `PMIx_Put`. Thus, applications should never use a defined “PMIX” attribute as the key in a call to `PMIx_Put`.

5.2.1.1 Scope of Put Data

PMIx v1.0

The `pmix_scope_t` structure is a `uint8_t` type that defines the availability of data passed to `PMIx_Put`. The following constants can be used to set a variable of the type `pmix_scope_t`. All definitions were introduced in version 1 of the standard unless otherwise marked.

Specific implementations may support different scope values, but all implementations must support at least `PMIX_GLOBAL`. If a specified scope value is not supported, then the `PMIx_Put` call must return `PMIX_ERR_NOT_SUPPORTED`.

`PMIX_SCOPE_UNDEF 0` Undefined scope.

`PMIX_LOCAL 1` The data is intended only for other application processes on the same node. Data marked in this way will not be included in data packages sent to remote requesters - i.e., it is only available to processes on the local node.

`PMIX_REMOTE 2` The data is intended solely for applications processes on remote nodes.

Data marked in this way will not be shared with other processes on the same node - i.e., it is only available to processes on remote nodes.

`PMIX_GLOBAL 3` The data is to be shared with all other requesting processes, regardless of location.

`PMIX_INTERNAL 4` The data is intended solely for this process and is not shared with other processes.

PMIx v2.0

1 5.2.2 PMIx_Store_internal

2 Summary

3 Store some data locally for retrieval by other areas of the process.

4 Format

PMIx v1.0

C

5 `pmix_status_t`

```
6 PMIx_Store_internal(const pmix_proc_t *proc,  
7                    const pmix_key_t key,  
8                    pmix_value_t *val);
```

C

9 **IN** `proc`

process reference (handle)

11 **IN** `key`

key to retrieve (string)

13 **IN** `val`

Value to store (handle)

15 Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- 16 • [PMIX_ERR_BAD_PARAM](#) indicating a reserved key is provided in the `key` argument.

17 If none of the above return codes are appropriate, then an implementation must return either a
18 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

19 Description

20 Store some data locally for retrieval by other areas of the process. This is data that has only internal
21 scope - it will never be posted externally. Typically used to cache data obtained by means outside of
22 PMIx so that it can be accessed by various areas of the process.

23 5.2.3 PMIx_Commit

24 Summary

25 Make available to other processes all key-value pairs previously staged via [PMIx_Put](#).

26 Format

PMIx v1.0

C

```
27 pmix_status_t PMIx_Commit(void);
```

C

28 Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Description

PMIx implementations may choose to locally cache non-reserved keys prior to submitting them for distribution. `PMIx_Commit` initiates the operation of making available previously staged key-value pairs to other processes. Depending on the implementation, this may involve transmitting the entire collection of data posted by the process to a server. `PMIx_Commit` is an asynchronous operation that will immediately return to the caller while the data is staged in the background.

Advice to users

Users are advised to always include the call to `PMIx_Commit` in case the local implementation requires it. Note that posted data will not necessarily be circulated during `PMIx_Commit`. Availability of the data by other processes upon completion of `PMIx_Commit` therefore still relies upon the exchange mechanisms described at the beginning of this chapter.

5.3 Retrieval rules for non-reserved keys

Since non-reserved keys cannot, by definition, have been provided by the host environment, their retrieval follows significantly different rules than those defined for reserved keys (as detailed in Section 6). `PMIx_Get` for a non-reserved key will obey the following precedence search:

1. If the `PMIX_GET_REFRESH_CACHE` attribute is given, then the request is first forwarded to the local PMIx server which will then update the client's cache. Note that this may not, depending upon implementation details, result in any action.
2. Check the local PMIx client cache for the requested key - if not found and either the `PMIX_OPTIONAL` or `PMIX_GET_REFRESH_CACHE` attribute was given, the search will stop at this point and return the `PMIX_ERR_NOT_FOUND` status.
3. Request the information from the local PMIx server. The server will check its cache for the specified key. If the value still isn't found and the `PMIX_IMMEDIATE` attribute was given, then the library shall return the `PMIX_ERR_NOT_FOUND` error constant to the requester. Otherwise, the PMIx server library will take one of the following actions:
 - If the target process has a rank of `PMIX_RANK_UNDEF`, then this indicates that the key being requested is globally unique and *not* associated with a specific process. In this case, the server shall hold the request until either the data appears at the server or, if given, the `PMIX_TIMEOUT` is reached. In the latter case, the server will return the `PMIX_ERR_TIMEOUT` status. Note that the server may, depending on PMIx implementation, never respond if the caller failed to specify a `PMIX_TIMEOUT` and the requested key fails to arrive at the server.

- If the target process is *local* (i.e., attached to the same PMIx server), then the server will hold the request until either the target process provides the data or, if given, the `PMIX_TIMEOUT` is reached. In the latter case, the server will return the `PMIX_ERR_TIMEOUT` status. Note that data which is posted via `PMIx_Put` but not staged with `PMIx_Commit` may, depending upon implementation, never appear at the server.
- If the target process is *remote* (i.e., not attached to the same PMIx server), the server will either:
 - If the host has provided the `pmix_server_dmodex_req_fn_t` module function interface, then the server shall pass the request to its host for servicing. The host is responsible for determining the location of the target process and passing the request to the PMIx server at that location.

When the remote data request is received, the target PMIx server will check its cache for the specified key. If the key is not present, the request shall be held until either the target process provides the data or, if given, the `PMIX_TIMEOUT` is reached. In the latter case, the server will return the `PMIX_ERR_TIMEOUT` status. The host shall convey the result back to the originating PMIx server, which will reply to the requesting client with the result of the request when the host provides it.

Note that the target server may, depending on PMIx implementation, never respond if the caller failed to specify a `PMIX_TIMEOUT` and the target process fails to post the requested key.

- if the host does not support the `pmix_server_dmodex_req_fn_t` interface, then the server will immediately respond to the client with the `PMIX_ERR_NOT_FOUND` status

Advice to PMIx library implementers

While there is no requirement that all PMIx implementations follow the client-server paradigm used in the above description, implementers are required to provide behaviors consistent with the described search pattern.

Advice to users

Users are advised to always specify the `PMIX_TIMEOUT` value when retrieving non-reserved keys to avoid potential deadlocks should the specified key not become available.

5.4 PMIx_Get

Summary

Retrieve a key/value pair from the client's namespace.

Format

C

```
pmix_status_t
PMIx_Get(const pmix_proc_t *proc, const pmix_key_t key,
         const pmix_info_t info[], size_t ninfo,
         pmix_value_t **val);
```

C

IN `proc`

Process identifier - a **NULL** value may be used in place of the caller's ID (handle)

IN `key`

Key to retrieve ([pmix_key_t](#))

IN `info`

Array of info structures (array of handles)

IN `ninfo`

Number of elements in the *info* array (integer)

OUT `val`

value (handle)

A successful return indicates that the requested data has been returned in the manner requested (.e.g., in a provided static memory location).

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- **PMIX_ERR_BAD_PARAM** A bad parameter was passed to the function call - e.g., the request included the **PMIX_GET_STATIC_VALUES** directive, but the provided storage location was **NULL**
- **PMIX_ERR_EXISTS_OUTSIDE_SCOPE** The requested key exists, but was posted in a *scope* (see Section 5.2.1.1) that does not include the requester.
- **PMIX_ERR_NOT_FOUND** The requested data was not available.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_OPTIONAL "pmix.optional" (bool)

Look only in the client's local data store for the requested value - do not request data from the PMIx server if not found.

PMIX_IMMEDIATE "pmix.immediate" (bool)

Specified operation should immediately return an error from the PMIx server if the requested data cannot be found - do not request it from the host RM.

PMIX_DATA_SCOPE "pmix.scope" ([pmix_scope_t](#))

1 Scope of the data to be searched in a **PMIx_Get** call.

2 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)

3 Return information regarding the session realm of the target process.

4 **PMIX_JOB_INFO** "pmix.job.info" (bool)

5 Return information regarding the job realm corresponding to the namespace in the target
6 process' identifier.

7 **PMIX_APP_INFO** "pmix.app.info" (bool)

8 Return information regarding the application realm to which the target process belongs - the
9 namespace of the target process serves to identify the job containing the target application. If
10 information about an application other than the one containing the target process is desired,
11 then the attribute array must contain a **PMIX_APPNUM** attribute identifying the desired
12 target application. This is useful in cases where there are multiple applications and the
13 mapping of processes to applications is unclear.

14 **PMIX_NODE_INFO** "pmix.node.info" (bool)

15 Return information from the node realm regarding the node upon which the specified
16 process is executing. If information about a node other than the one containing the specified
17 process is desired, then the attribute array must also contain either the **PMIX_NODEID** or
18 **PMIX_HOSTNAME** attribute identifying the desired target. This is useful for requesting
19 information about a specific node even if the identity of processes running on that node are
20 not known.

21 **PMIX_GET_STATIC_VALUES** "pmix.get.static" (bool)

22 Request that the data be returned in the provided storage location. The caller is responsible
23 for destructing the **pmix_value_t** using the **PMIX_VALUE_DESTRUCT** macro when
24 done.

25 **PMIX_GET_POINTER_VALUES** "pmix.get.pntrs" (bool)

26 Request that any pointers in the returned value point directly to values in the key-value store.
27 The user *must not* release any returned data pointers.

28 **PMIX_GET_REFRESH_CACHE** "pmix.get.refresh" (bool)

29 When retrieving data for a remote process, refresh the existing local data cache for the
30 process in case new values have been put and committed by the process since the last refresh.
31 Local process information is assumed to be automatically updated upon posting by the
32 process. A **NULL** key will cause all values associated with the process to be refreshed -
33 otherwise, only the indicated key will be updated. A process rank of
34 **PMIX_RANK_WILDCARD** can be used to update job-related information in dynamic
35 environments. The user is responsible for subsequently updating refreshed values they may
36 have cached in their own local memory.

Optional Attributes

The following attributes are optional for host environments:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Retrieve information for the specified *key* associated with the process identified in the given **pmix_proc_t**. See Chapters 6 and 5.1 for details on rules governing retrieval of information. Information will be returned according to provided directives:

- In the absence of any directive, the returned **pmix_value_t** shall be an allocated memory object. The caller is responsible for releasing the object when done.
- If **PMIX_GET_POINTER_VALUES** is given, then the function shall return a pointer to a **pmix_value_t** in the PMIx library's memory that contains the requested information.
- If **PMIX_GET_STATIC_VALUES** is given, then the function shall return the information in the provided **pmix_value_t** pointer. In this case, the caller must provide storage for the structure and pass the pointer to that storage in the *val* parameter. If the implementation cannot return a static value, then the call to **PMIx_Get** must return the **PMIX_ERR_NOT_SUPPORTED** status.

Retrieve information for the specified *key* associated with the process identified in the given **pmix_proc_t**. See Chapters 6 and 5.1 for details on rules governing retrieval of information. Information will be returned according to provided directives:

- In the absence of any directive, the returned **pmix_value_t** shall be an allocated memory object. The caller is responsible for releasing the object when done.
- If **PMIX_GET_POINTER_VALUES** is given, then the function shall return a pointer to a **pmix_value_t** in the PMIx library's memory that contains the requested information.
- If **PMIX_GET_STATIC_VALUES** is given, then the function shall return the information in the provided **pmix_value_t** pointer. In this case, the caller must provide storage for the structure and pass the pointer to that storage in the *val* parameter.

This is a blocking operation - the caller will block until the retrieval rules of Section 6.2 or 5.3 are met.

The *info* array is used to pass user directives regarding the get operation.

5.4.1 PMIx_Get_nb

Summary

Nonblocking **PMIx_Get** operation.

Format

C

```
pmix_status_t
PMIx_Get_nb(const pmix_proc_t *proc, const char key[],
            const pmix_info_t info[], size_t ninfo,
            pmix_value_cbfunc_t cbfunc, void *cbdata);
```

C

IN `proc`

Process identifier - a **NULL** value may be used in place of the caller's ID (handle)

IN `key`

Key to retrieve (string)

IN `info`

Array of info structures (array of handles)

IN `ninfo`

Number of elements in the *info* array (integer)

IN `cbfunc`

Callback function (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

If executed, the status returned in the provided callback function will be one of the following constants:

- **PMIX_SUCCESS** The requested data has been returned.
- **PMIX_ERR_EXISTS_OUTSIDE_SCOPE** The requested key exists, but was posted in a *scope* (see Section 5.2.1.1) that does not include the requester.
- **PMIX_ERR_NOT_FOUND** The requested data was not available.
- a non-zero PMIx error constant indicating a reason for the request's failure.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_OPTIONAL "pmix.optional" (bool)

Look only in the client's local data store for the requested value - do not request data from the PMIx server if not found.

PMIX_IMMEDIATE "pmix.immediate" (bool)

Specified operation should immediately return an error from the PMIx server if the requested data cannot be found - do not request it from the host RM.

1 **PMIX_DATA_SCOPE** "pmix.scope" (pmix_scope_t)
 2 Scope of the data to be searched in a **PMIx_Get** call.

3 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)
 4 Return information regarding the session realm of the target process.

5 **PMIX_JOB_INFO** "pmix.job.info" (bool)
 6 Return information regarding the job realm corresponding to the namespace in the target
 7 process' identifier.

8 **PMIX_APP_INFO** "pmix.app.info" (bool)
 9 Return information regarding the application realm to which the target process belongs - the
 10 namespace of the target process serves to identify the job containing the target application. If
 11 information about an application other than the one containing the target process is desired,
 12 then the attribute array must contain a **PMIX_APPNUM** attribute identifying the desired
 13 target application. This is useful in cases where there are multiple applications and the
 14 mapping of processes to applications is unclear.

15 **PMIX_NODE_INFO** "pmix.node.info" (bool)
 16 Return information from the node realm regarding the node upon which the specified
 17 process is executing. If information about a node other than the one containing the specified
 18 process is desired, then the attribute array must also contain either the **PMIX_NODEID** or
 19 **PMIX_HOSTNAME** attribute identifying the desired target. This is useful for requesting
 20 information about a specific node even if the identity of processes running on that node are
 21 not known.

22 **PMIX_GET_POINTER_VALUES** "pmix.get.pntrs" (bool)
 23 Request that any pointers in the returned value point directly to values in the key-value store.
 24 The user *must not* release any returned data pointers.

25 **PMIX_GET_REFRESH_CACHE** "pmix.get.refresh" (bool)
 26 When retrieving data for a remote process, refresh the existing local data cache for the
 27 process in case new values have been put and committed by the process since the last refresh.
 28 Local process information is assumed to be automatically updated upon posting by the
 29 process. A **NULL** key will cause all values associated with the process to be refreshed -
 30 otherwise, only the indicated key will be updated. A process rank of
 31 **PMIX_RANK_WILDCARD** can be used to update job-related information in dynamic
 32 environments. The user is responsible for subsequently updating refreshed values they may
 33 have cached in their own local memory.

34

35 The following attributes are required for host environments that support this operation:

36 **PMIX_WAIT** "pmix.wait" (int)
 37 Caller requests that the PMIx server wait until at least the specified number of values are
 38 found (a value of zero indicates *all* and is the default).

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

The callback function will be executed once the retrieval rules of Chapters 6 or 5.3 are met. See **PMIx_Get** for a full description. Note that the non-blocking form of this function cannot support the **PMIX_GET_STATIC_VALUES** attribute as the user cannot pass in the required pointer to storage for the result.

5.4.2 Retrieval attributes

The following attributes are defined for use by retrieval APIs:

PMIX_OPTIONAL "pmix.optional" (bool)

Look only in the client's local data store for the requested value - do not request data from the PMIx server if not found.

PMIX_IMMEDIATE "pmix.immediate" (bool)

Specified operation should immediately return an error from the PMIx server if the requested data cannot be found - do not request it from the host RM.

PMIX_GET_POINTER_VALUES "pmix.get.pntrs" (bool)

Request that any pointers in the returned value point directly to values in the key-value store. The user *must not* release any returned data pointers.

PMIX_GET_STATIC_VALUES "pmix.get.static" (bool)

Request that the data be returned in the provided storage location. The caller is responsible for destructing the **pmix_value_t** using the **PMIX_VALUE_DESTRUCT** macro when done.

PMIX_GET_REFRESH_CACHE "pmix.get.refresh" (bool)

When retrieving data for a remote process, refresh the existing local data cache for the process in case new values have been put and committed by the process since the last refresh.

Local process information is assumed to be automatically updated upon posting by the process. A **NULL** key will cause all values associated with the process to be refreshed - otherwise, only the indicated key will be updated. A process rank of

PMIX_RANK_WILDCARD can be used to update job-related information in dynamic environments. The user is responsible for subsequently updating refreshed values they may have cached in their own local memory.

PMIX_DATA_SCOPE "pmix.scope" (**pmix_scope_t**)

1 Scope of the data to be searched in a `PMIx_Get` call.

2 **PMIX_TIMEOUT** "`pmix.timeout`" (`int`)

3 Time in seconds before the specified operation should time out (zero indicating infinite) and
4 return the `PMIX_ERR_TIMEOUT` error. Care should be taken to avoid race conditions
5 caused by multiple layers (client, server, and host) simultaneously timing the operation.

6 **PMIX_WAIT** "`pmix.wait`" (`int`)

7 Caller requests that the PMIx server wait until at least the specified number of values are
8 found (a value of zero indicates *all* and is the default).

CHAPTER 6

Reserved Keys

1 *Reserved* keys are keys whose string representation begin with a prefix of "**pmix**". By definition,
2 reserved keys are provided by the host environment and the PMIx server, and are required to be
3 available at client start of execution. PMIx clients and tools are therefore prohibited from posting
4 reserved keys.

5 Host environments may opt to define non-standardized reserved keys. All reserved keys, whether
6 standardized or non-standardized, follow the same retrieval rules. Users are advised to check both
7 the local PMIx implementation and host environment documentation for a list of any
8 non-standardized reserved keys they must avoid, and to learn of any non-standard keys that may
9 require special handling.

10 6.1 Data realms

11 PMIx information spans a wide range of sources. In some cases, there are multiple overlapping
12 sources for the same type of data - e.g., the session, job, and application can each provide
13 information on the number of nodes involved in their respective area. In order to resolve the
14 ambiguity, a *data realm* is used to identify the scope to which the referenced data applies. Thus, a
15 reference to an attribute that isn't specific to a realm (e.g., the **PMIX_NUM_NODES** attribute) must
16 be accompanied by a corresponding attribute identifying the realm to which the request pertains if
17 it differs from the default.

18 PMIx defines five *data realms* to resolve the ambiguities, as captured in the following attributes
19 used in **PMIx_Get** for retrieving information from each of the realms:

20 **PMIX_SESSION_INFO** "**pmix.ssn.info**" (**bool**)

21 Return information regarding the session realm of the target process.

22 **PMIX_JOB_INFO** "**pmix.job.info**" (**bool**)

23 Return information regarding the job realm corresponding to the namespace in the target
24 process' identifier.

25 **PMIX_APP_INFO** "**pmix.app.info**" (**bool**)

26 Return information regarding the application realm to which the target process belongs - the
27 namespace of the target process serves to identify the job containing the target application. If
28 information about an application other than the one containing the target process is desired,
29 then the attribute array must contain a **PMIX_APPNUM** attribute identifying the desired
30 target application. This is useful in cases where there are multiple applications and the
31 mapping of processes to applications is unclear.

1 **PMIX_PROC_INFO** "pmix.proc.info" (bool)

2 Return information regarding the target process. This attribute is technically not required as
3 the **PMIx_Get** API specifically identifies the target process in its parameters. However, it is
4 included here for completeness.

5 **PMIX_NODE_INFO** "pmix.node.info" (bool)

6 Return information from the node realm regarding the node upon which the specified
7 process is executing. If information about a node other than the one containing the specified
8 process is desired, then the attribute array must also contain either the **PMIX_NODEID** or
9 **PMIX_HOSTNAME** attribute identifying the desired target. This is useful for requesting
10 information about a specific node even if the identity of processes running on that node are
11 not known.

Advice to users

12 If information about a session other than the one containing the requesting process is desired, then
13 the attribute array must contain a **PMIX_SESSION_ID** attribute identifying the desired target
14 session. This is required as many environments only guarantee unique namespaces within a
15 session, and not across sessions.

16 Determining the target within a realm varies between realms and is explained in detail in the realm
17 descriptions below. Note that several attributes can be either queried as a key or set as an attribute
18 to specify the target within a realm. The attributes **PMIX_SESSION_ID**, **PMIX_NAMESPACE** and
19 **PMIX_APPNUM** can be used in both ways.

6.1.1 Session realm attributes

20 If information about a session other than the one containing the requesting process is desired, then
21 the *info* array passed to **PMIx_Get** must contain a **PMIX_SESSION_ID** attribute identifying the
22 desired target session. This is required as many environments only guarantee unique namespaces
23 within a session, and not across sessions.
24

25 Note that the *proc* argument of **PMIx_Get** is ignored when referencing session-related
26 information.

27 The following keys, by default, request session-level information. They will return information
28 about the caller's session unless a **PMIX_SESSION_ID** attribute is specified in the *info* array
29 passed to **PMIx_Get**:

30 **PMIX_CLUSTER_ID** "pmix.clid" (char*)

31 A string name for the cluster this allocation is on.

32 **PMIX_UNIV_SIZE** "pmix.univ.size" (uint32_t)

33 Maximum number of process that can be simultaneously executing in a session. Note that
34 this attribute is equivalent to the **PMIX_MAX_PROCS** attribute for the *session* realm - it is
35 included in the PMIx Standard for historical reasons.

36 **PMIX_TMPDIR** "pmix.tmpdir" (char*)

1 Full path to the top-level temporary directory assigned to the session.
2 **PMIX_TDIR_RMCLEAN** "pmix.tdir.rmclean" (bool)
3 The Resource Manager will remove any directories or files it creates in **PMIX_TMPDIR**.
4 **PMIX_HOSTNAME_KEEP_FQDN** "pmix.fqdn" (bool)
5 Fully Qualified Domain Names (FQDNs) are being retained by the PMIx library.
6 **PMIX_RM_NAME** "pmix.rm.name" (char*)
7 String name of the RM.
8 **PMIX_RM_VERSION** "pmix.rm.version" (char*)
9 RM version string.

10 The following session-related keys default to the realms described in their descriptions but can be
11 retrieved from the session realm by setting the **PMIX_SESSION_INFO** attribute in the *info* array
12 passed to **PMIx_Get**:

13 **PMIX_ALLOCATED_NODELIST** "pmix.alist" (char*)
14 Comma-delimited list or regular expression of all nodes in the specified realm regardless of
15 whether or not they currently host processes. Defaults to the *job* realm.
16 **PMIX_NUM_ALLOCATED_NODES** "pmix.num.anodes" (uint32_t)
17 Number of nodes in the specified realm regardless of whether or not they currently host
18 processes. Defaults to the *job* realm.
19 **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)
20 Maximum number of processes that can be simultaneously executed in the specified realm.
21 Typically, this is a constraint imposed by a scheduler or by user settings in a hostfile or other
22 resource description. Defaults to the *job* realm.
23 **PMIX_NODE_LIST** "pmix.nlist" (char*)
24 Comma-delimited list of nodes currently hosting processes in the specified realm. Defaults
25 to the *job* realm.
26 **PMIX_NUM_SLOTS** "pmix.num.slots" (uint32_t)
27 Maximum number of processes that can simultaneously be executing in the specified realm.
28 Note that this attribute is the equivalent to **PMIX_MAX_PROCS** - it is included in the PMIx
29 Standard for historical reasons. Defaults to the *job* realm.
30 **PMIX_NUM_NODES** "pmix.num.nodes" (uint32_t)
31 Number of nodes currently hosting processes in the specified realm. Defaults to the *job*
32 realm.
33 **PMIX_NODE_MAP** "pmix.nmap" (char*)
34 Regular expression of nodes currently hosting processes in the specified realm - see [17.2.3.2](#)
35 for an explanation of its generation. Defaults to the *job* realm.
36 **PMIX_NODE_MAP_RAW** "pmix.nmap.raw" (char*)
37 Comma-delimited list of nodes containing procs within the specified realm. Defaults to the
38 *job* realm.
39 **PMIX_PROC_MAP** "pmix.pmap" (char*)
40 Regular expression describing processes on each node in the specified realm - see [17.2.3.2](#)
41 for an explanation of its generation. Defaults to the *job* realm.
42 **PMIX_PROC_MAP_RAW** "pmix.pmap.raw" (char*)

1 Semi-colon delimited list of strings, each string containing a comma-delimited list of ranks
2 on the corresponding node within the specified realm. Defaults to the *job* realm.

3 **PMIX_ANL_MAP** "pmix.anlmap" (char*)

4 Process map equivalent to **PMIX_PROC_MAP** expressed in Argonne National Laboratory's
5 PMI-1/PMI-2 notation. Defaults to the *job* realm.

6 6.1.2 Job realm attributes

7 Job-related information can be retrieved by requesting a key which defaults to the job realm or by
8 including the **PMIX_JOB_INFO** attribute in the *info* array passed to **PMIx_Get**. For job-related
9 keys the target job is specified by setting the namespace of the target job in the *proc* argument and
10 specifying a rank of **PMIX_RANK_WILDCARD** in the *proc* argument passed to **PMIx_Get**.

11 If information is requested about a namespace in a session other than the one containing the
12 requesting process, then the *info* array must contain a **PMIX_SESSION_ID** attribute identifying
13 the desired target session. This is required as many environments only guarantee unique
14 namespaces within a session, and not across sessions.

15 The following keys, by default, request job-level information: They will return information about
16 the job indicated in *proc*:

17 **PMIX_JOBID** "pmix.jobid" (char*)

18 Job identifier assigned by the scheduler to the specified job - may be identical to the
19 namespace, but is often a numerical value expressed as a string (e.g., "12345.3").

20 **PMIX_NPROC_OFFSET** "pmix.offset" (pmix_rank_t)

21 Starting global rank of the specified job. The returned value is the same as the value of
22 **PMIX_GLOBAL_RANK** of rank 0 of the specified job.

23 **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)

24 Maximum number of processes that can be simultaneously executed in the specified realm.
25 Typically, this is a constraint imposed by a scheduler or by user settings in a hostfile or other
26 resource description. Defaults to the *job* realm. In this context, this is the maximum number
27 of processes that can be simultaneously executed in the specified job, which may be a subset
28 of the number allocated to the overall session.

29 **PMIX_NUM_SLOTS** "pmix.num.slots" (uint32_t)

30 Maximum number of processes that can simultaneously be executing in the specified realm.
31 Note that this attribute is the equivalent to **PMIX_MAX_PROCS** - it is included in the PMIx
32 Standard for historical reasons. Defaults to the *job* realm. In this context, this is the
33 maximum number of process that can be simultaneously executing within the specified job,
34 which may be a subset of the number allocated to the overall session. Jobs may reserve a
35 subset of their assigned maximum processes for dynamic operations such as **PMIx_Spawn**.

36 **PMIX_NUM_NODES** "pmix.num.nodes" (uint32_t)

37 Number of nodes currently hosting processes in the specified realm. Defaults to the *job*
38 realm. In this context, this is the number of nodes currently hosting processes in the
39 specified job, which may be a subset of the nodes allocated to the overall session. Jobs may

1 reserve a subset of their assigned nodes for dynamic operations such as **PMIx_Spawn** - i.e.,
2 not all nodes may have executing processes from this job at a given point in time.

3 **PMIX_NODE_MAP** "pmix.nmap" (char*)

4 Regular expression of nodes currently hosting processes in the specified realm - see 17.2.3.2
5 for an explanation of its generation. Defaults to the *job* realm. In this context, this is the
6 regular expression of nodes currently hosting processes in the specified job.

7 **PMIX_NODE_LIST** "pmix.nlist" (char*)

8 Comma-delimited list of nodes currently hosting processes in the specified realm. Defaults
9 to the *job* realm. In this context, this is the comma-delimited list of nodes currently hosting
10 processes in the specified job.

11 **PMIX_PROC_MAP** "pmix.pmap" (char*)

12 Regular expression describing processes on each node in the specified realm - see 17.2.3.2
13 for an explanation of its generation. Defaults to the *job* realm. In this context, this is the
14 regular expression describing processes on each node in the specified job.

15 **PMIX_ANL_MAP** "pmix.anlmap" (char*)

16 Process map equivalent to **PMIX_PROC_MAP** expressed in Argonne National Laboratory's
17 PMI-1/PMI-2 notation. Defaults to the *job* realm. In this context, this is the process
18 mapping in Argonne National Laboratory's PMI-1/PMI-2 notation of the processes in the
19 specified job.

20 **PMIX_CMD_LINE** "pmix.cmd.line" (char*)

21 Command line used to execute the specified job (e.g., "mpirun -n 2 --map-by foo ./myapp : -n
22 4 ./myapp2"). If the job was created by a call to **PMIx_Spawn**, the string is an in-order
23 concatenation of the values of **PMIX_APP_ARGV** for each application in the job using the
24 character ':' as a separator.

25 **PMIX_NSDIR** "pmix.nmdir" (char*)

26 Full path to the temporary directory assigned to the specified job, under **PMIX_TMPDIR**.

27 **PMIX_JOB_SIZE** "pmix.job.size" (uint32_t)

28 Total number of processes in the specified job across all contained applications. Note that
29 this value can be different from **PMIX_MAX_PROCS**. For example, users may choose to
30 subdivide an allocation (running several jobs in parallel within it), and dynamic
31 programming models may support adding and removing processes from a running *job*
32 on-the-fly. In the latter case, PMIx events may be used to notify processes within the job that
33 the job size has changed.

34 **PMIX_JOB_NUM_APPS** "pmix.job.napps" (uint32_t)

35 Number of applications in the specified job.

36 **PMIX_LOCAL_PEERS** "pmix.lpeers" (char*)

37 Comma-delimited list of ranks that are executing on the local node within the specified
38 namespace – shortcut for **PMIx_Resolve_peers** for the local node.

39 **PMIX_LOCALLDR** "pmix.llldr" (pmix_rank_t)

40 Lowest rank within the specified job on the node (defaults to current node in absence of
41 **PMIX_HOSTNAME** or **PMIX_NODEID** qualifier).

1 **PMIX_LOCAL_CPUSSETS** "pmix.lcpus" (**pmix_data_array_t**)
 2 A **pmix_data_array_t** array of string representations of the Processing Unit (PU)
 3 binding bitmaps applied to each local *peer* on the caller's node upon launch. Each string
 4 shall begin with the name of the library that generated it (e.g., "hwloc") followed by a colon
 5 and the bitmap string itself. The array shall be in the same order as the processes returned by
 6 **PMIX_LOCAL_PEERS** for that namespace.
 7 **PMIX_LOCAL_SIZE** "pmix.local.size" (**uint32_t**)
 8 Number of processes in the specified job or application on the caller's node. Defaults to job
 9 unless the **PMIX_APP_INFO** and the **PMIX_APPNUM** qualifiers are given.

10 6.1.3 Application realm attributes

11 Application-related information can be retrieved by requesting a key which defaults to the
 12 application realm or by including the **PMIX_APP_INFO** attribute in the *info* array passed to
 13 **PMIx_Get**. If the **PMIX_APPNUM** qualifier is given, then the query shall return the corresponding
 14 value for the given application within the namespace specified in the *proc* argument of the query (a
 15 **NULL** value for the *proc* argument equates to the namespace of the caller). If the **PMIX_APPNUM**
 16 qualifier is not included, then the retrieval shall default to the application containing the process
 17 specified by *proc*. If the rank specified in *proc* is **PMIX_RANK_WILDCARD**, then the application
 18 number shall default to that of the calling process if the namespace is its own job, or a value of zero
 19 if the namespace is that of a different job.

20 The following keys, by default, request application-level information. They will return information
 21 about the application indicated in *proc*:

22 **PMIX_APPLDR** "pmix.aldr" (**pmix_rank_t**)
 23 Lowest rank in the specified application.
 24 **PMIX_APP_SIZE** "pmix.app.size" (**uint32_t**)
 25 Number of processes in the specified application, regardless of their execution state - i.e.,
 26 this number may include processes that either failed to start or have already terminated.
 27 **PMIX_APP_ARGV** "pmix.app.argv" (**char***)
 28 Consolidated argv passed to the spawn command for the given application (e.g., "./myapp
 29 arg1 arg2 arg3").
 30 **PMIX_APP_MAP_TYPE** "pmix.apmap.type" (**char***)
 31 Type of mapping used to layout the application (e.g., **cyclic**).
 32 **PMIX_APP_MAP_REGEX** "pmix.apmap.regex" (**char***)
 33 Regular expression describing the result of the process mapping.

34 The following application-related keys default to the realms described in their descriptions but can
 35 be retrieved from the application realm by setting the **PMIX_APP_INFO** attribute in the *info* array
 36 passed to **PMIx_Get**:

37 **PMIX_NUM_NODES** "pmix.num.nodes" (**uint32_t**)
 38 Number of nodes currently hosting processes in the specified realm. Defaults to the *job*
 39 realm. In this context, this is the number of nodes currently hosting processes in the
 40 specified application, which may be a subset of the nodes allocated to the overall session.

1 **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)

2 Maximum number of processes that can be simultaneously executed in the specified realm.
3 Typically, this is a constraint imposed by a scheduler or by user settings in a hostfile or other
4 resource description. Defaults to the *job* realm. In this context, this is the maximum number
5 of processes that can be executed in the specified application, which may be a subset of the
6 number allocated to the overall session and job.

7 **PMIX_NUM_SLOTS** "pmix.num.slots" (uint32_t)

8 Maximum number of processes that can simultaneously be executing in the specified realm.
9 Note that this attribute is the equivalent to **PMIX_MAX_PROCS** - it is included in the PMIx
10 Standard for historical reasons. Defaults to the *job* realm. In this context, this is the number
11 of slots assigned to the specified application, which may be a subset of the slots allocated to
12 the overall session and job.

13 **PMIX_NODE_MAP** "pmix.nmap" (char*)

14 Regular expression of nodes currently hosting processes in the specified realm - see 17.2.3.2
15 for an explanation of its generation. Defaults to the *job* realm. In this context, this is the
16 regular expression of nodes currently hosting processes in the specified application.

17 **PMIX_NODE_LIST** "pmix.nlist" (char*)

18 Comma-delimited list of nodes currently hosting processes in the specified realm. Defaults
19 to the *job* realm. In this context, this is the comma-delimited list of nodes currently hosting
20 processes in the specified application.

21 **PMIX_PROC_MAP** "pmix.pmap" (char*)

22 Regular expression describing processes on each node in the specified realm - see 17.2.3.2
23 for an explanation of its generation. Defaults to the *job* realm. In this context, this is the
24 regular expression describing processes on each node in the specified application.

25 6.1.4 Process realm attributes

26 Process-related information can be retrieved by requesting a key which defaults to the process
27 realm or by including the **PMIX_PROC_INFO** attribute in the *info* array passed to **PMIx_Get**.
28 The target process is specified by the namespace and rank of the *proc* argument to **PMIx_Get**. For
29 process-related keys (other than **PMIX_PROCID** and **PMIX_NAMESPACE**) the target process is
30 specified by setting the namespace and rank of the target process in the *proc* argument passed to
31 **PMIx_Get**. If information is requested about a process in a session other than the one containing
32 the requesting process, then an attribute identifying the target session must be provided. This is
33 required as many environments only guarantee unique namespaces within a session, and not across
34 sessions.

35 The following keys, by default, request process-level information: They will return information
36 about the process indicated in *proc*:

37 **PMIX_APPNUM** "pmix.appnum" (uint32_t)

38 The application number within the job in which the specified process is a member.

1 **PMIX_RANK** "pmix.rank" (pmix_rank_t)
 2 Process rank within the job, starting from zero.

3 **PMIX_NAMESPACE** "pmix.namespace" (char*)
 4 Namespace of the job - may be a numerical value expressed as a string, but is often an
 5 alphanumeric string carrying information solely of use to the system. Required to be unique
 6 within the scope of the host environment. One cannot retrieve the namespace of an arbitrary
 7 process since that would require already knowing the namespace of that process. However, a
 8 process' own namespace can be retrieved by passing a NULL value of *proc* to **PMIx_Get**.

9 **PMIX_SESSION_ID** "pmix.session.id" (uint32_t)
 10 Session identifier assigned by the scheduler.

11 **PMIX_GLOBAL_RANK** "pmix.grank" (pmix_rank_t)
 12 Rank of the specified process spanning across all jobs in this session, starting with zero.
 13 Note that no ordering of the jobs is implied when computing this value. As jobs can start and
 14 end at random times, this is defined as a continually growing number - i.e., it is not
 15 dynamically adjusted as individual jobs and processes are started or terminated.

16 **PMIX_APP_RANK** "pmix.apprank" (pmix_rank_t)
 17 Rank of the specified process within its application.

18 **PMIX_PARENT_ID** "pmix.parent" (pmix_proc_t)
 19 Process identifier of the parent process of the specified process - typically used to identify
 20 the application process that caused the job containing the specified process to be spawned
 21 (e.g., the process that called **PMIx_Spawn**). This attribute is only provided for a process if
 22 it was created by a call to **PMIx_Spawn** or **PMIx_Spawn_nb**.

23 **PMIX_EXIT_CODE** "pmix.exit.code" (int)
 24 Exit code returned when the specified process terminated.

25 **PMIX_PROCID** "pmix.procid" (pmix_proc_t)
 26 The caller's process identifier. The value returned is identical to what **PMIx_Init** or
 27 **PMIx_tool_init** provides. The process identifier in the **PMIx_Get** call is ignored when
 28 requesting this key.

29 **PMIX_LOCAL_RANK** "pmix.lrank" (uint16_t)
 30 Rank of the specified process on its node - refers to the numerical location (starting from
 31 zero) of the process on its node when counting only those processes from the same job that
 32 share the node, ordered by their overall rank within that job.

33 **PMIX_NODE_RANK** "pmix.nrank" (uint16_t)
 34 Rank of the specified process on its node spanning all jobs- refers to the numerical location
 35 (starting from zero) of the process on its node when counting all processes (regardless of
 36 job) that share the node, ordered by their overall rank within the job. The value represents a
 37 snapshot in time when the specified process was started on its node and is not dynamically
 38 adjusted as processes from other jobs are started or terminated on the node.

39 **PMIX_PACKAGE_RANK** "pmix.pkgrank" (uint16_t)
 40 Rank of the specified process on the *package* where this process resides - refers to the
 41 numerical location (starting from zero) of the process on its package when counting only
 42 those processes from the same job that share the package, ordered by their overall rank

1 within that job. Note that processes that are not bound to PUs within a single specific
2 package cannot have a package rank.

3 **PMIX_PROC_PID** "pmix.ppid" (pid_t)

4 Operating system PID of specified process.

5 **PMIX_PROCDIR** "pmix.pdir" (char*)

6 Full path to the subdirectory under **PMIX_NSDIR** assigned to the specified process.

7 **PMIX_CPuset** "pmix.cpuset" (char*)

8 A string representation of the PU binding bitmap applied to the process upon launch. The
9 string shall begin with the name of the library that generated it (e.g., "hwloc") followed by a
10 colon and the bitmap string itself.

11 **PMIX_CPuset_BITMAP** "pmix.bitmap" (pmix_cpuset_t*)

12 Bitmap applied to the process upon launch.

13 **PMIX_CREDENTIAL** "pmix.cred" (char*)

14 Security credential assigned to the process.

15 **PMIX_SPawnED** "pmix.spawned" (bool)

16 **true** if this process resulted from a call to **PMIx_Spawn**. Lack of inclusion (i.e., a return
17 status of **PMIX_ERR_NOT_FOUND**) corresponds to a value of **false** for this attribute.

18 **PMIX_REINcARNATION** "pmix.reinc" (uint32_t)

19 Number of times this process has been re-instantiated - i.e, a value of zero indicates that the
20 process has never been restarted.

21 In addition, process-level information includes functional attributes directly associated with a
22 process - for example, the process-related fabric attributes included in Section 15.3 or the distance
23 attributes of Section 12.4.11.

24 6.1.5 Node realm keys

25 Node-related information can be retrieved by requesting a key which defaults to the node realm or
26 by including the **PMIX_NODE_INFO** attribute in the *info* array passed to **PMIx_Get**. The target
27 node defaults to the local node unless a different node is specified in the *info* array using either the
28 **PMIX_HOSTNAME** or **PMIX_NODEID**. Some node related keys are an exception to this rule and
29 are listed separately at the end of this section. These special keys can only target the local node and
30 also require that a namespace be specified using the *proc* argument to **PMIx_Get**.

31 The following keys, by default, request node-level information. They will return information about
32 either the local node or the node specified by **PMIX_HOSTNAME** or **PMIX_NODEID**:

33 **PMIX_HOSTNAME** "pmix.hname" (char*)

34 Name of the host, as returned by the **gethostname** utility or its equivalent.

35 **PMIX_HOSTNAME_ALIASES** "pmix.alias" (char*)

36 Comma-delimited list of names by which the target node is known.

37 **PMIX_NODEID** "pmix.nodeid" (uint32_t)

1 Node identifier expressed as the node's index (beginning at zero) in an array of nodes within
2 the active session. The value must be unique and directly correlate to the **PMIX_HOSTNAME**
3 of the node - i.e., users can interchangeably reference the same location using either the
4 **PMIX_HOSTNAME** or corresponding **PMIX_NODEID**.
5 **PMIX_NODE_SIZE** "pmix.node.size" (uint32_t)
6 Number of processes across all jobs that are executing upon the node.
7 **PMIX_AVAIL_PHYS_MEMORY** "pmix.pmem" (uint64_t)
8 Total available physical memory on a node.
9 **PMIX_LOCAL_PROCS** "pmix.lprocs" (pmix_proc_t array)
10 Array of **pmix_proc_t** of all processes executing on the local node – shortcut for
11 **PMIX_Resolve_peers** for the local node and a **NULL** namespace argument. The process
12 identifier is ignored for this attribute.
13 **Provisional** **PMIX_NODE_OVERSUBSCRIBED** "pmix.ndosub" (bool)
14 True if the number of processes from this job on this node exceeds the number of slots
15 allocated to it

16 In addition, node-level information includes functional attributes directly associated with a node -
17 for example, the node-related fabric attributes included in Section 15.3.

18 6.2 Retrieval rules for reserved keys

19 The retrieval rules for reserved keys are relatively simple as the keys, if provided by an
20 implementation, are required, by definition, to be available when the client begins execution.
21 Accordingly, **PMIX_Get** for a reserved key first checks the local PMIX Client cache (per the data
22 realm rules of the prior section) for the target key. If the information is not found, then the
23 **PMIX_ERR_NOT_FOUND** error constant is returned unless the target process belongs to a different
24 namespace from that of the requester.

25 In the case where the target and requester's namespaces differ, then the request is forwarded to the
26 local PMIX server. Upon receiving the request, the server shall check its data storage for the
27 specified namespace. If it already knows about this namespace, then it shall attempt to lookup the
28 specified key, returning the value if it is found or the **PMIX_ERR_NOT_FOUND** error constant.

29 If the server does not have a copy of the information for the specified namespace, then the server
30 shall take one of the following actions:

- 31 1. If the request included the **PMIX_IMMEDIATE** attribute, then the server will respond to the
32 client with the **PMIX_ERR_NOT_FOUND** status.
- 33 2. If the host has provided the Direct Business Card Exchange (DBCX) module function interface
34 (**pmix_server_dmodex_req_fn_t**), then the server shall pass the request to its host for
35 servicing. The host is responsible for identifying a source of information on the specified
36 namespace and retrieving it. The host is required to retrieve *all* of the information regarding the
37 target namespace and return it to the requesting server in anticipation of follow-on requests. If
38 the host cannot retrieve the namespace information, then it must respond with the

1 **PMIX_ERR_NOT_FOUND** error constant unless the **PMIX_TIMEOUT** is given and reached (in
2 which case, the host must respond with the **PMIX_ERR_TIMEOUT** constant).

3 Once the the PMIx server receives the namespace information, the server shall search it (again
4 adhering to the prior data realm rules) for the requested key, returning the value if it is found or
5 the **PMIX_ERR_NOT_FOUND** error constant.

- 6 3. If the host does not support the DBCX interface, then the server will respond to the client with
7 the **PMIX_ERR_NOT_FOUND** status

8 6.2.1 Accessing information: examples

9 This section provides examples illustrating methods for accessing information from the various
10 realms. The intent of the examples is not to provide comprehensive coding guidance, but rather to
11 further illustrate the use of **PMIx_Get** for obtaining information on a *session*, *job*, *application*,
12 *process*, and *node*.

13 6.2.1.1 Session-level information

14 The **PMIx_Get** API does not include an argument for specifying the *session* associated with the
15 information being requested. Thus, requests for keys that are not specifically for session-level
16 information must be accompanied by the **PMIX_SESSION_INFO** qualifier.

17 Example requests are shown below:

```
18 pmix_info_t info;  
19 pmix_value_t *value;  
20 pmix_status_t rc;  
21 pmix_proc_t myproc, wildcard;  
22  
23 /* initialize the client library */  
24 PMIx_Init(&myproc, NULL, 0);  
25  
26 /* get the #slots in our session */  
27 PMIX_PROC_LOAD(&wildcard, myproc.nspace, PMIX_RANK_WILDCARD);  
28 rc = PMIx_Get(&wildcard, PMIX_UNIV_SIZE, NULL, 0, &value);  
29  
30 /* get the #nodes in our session */  
31 PMIx_Info_load(&info, PMIX_SESSION_INFO, NULL, PMIX_BOOL);  
32 rc = PMIx_Get(&wildcard, PMIX_NUM_NODES, &info, 1, &value);
```

33 Information regarding a different session can be requested by adding the **PMIX_SESSION_ID**
34 attribute identifying the target session. In this case, the *proc* argument to **PMIx_Get** will be
35 ignored:

C

```
1 pmix_info_t info[2];
2 pmix_value_t *value;
3 pmix_status_t rc;
4 pmix_proc_t myproc;
5 uint32_t sid;
6
7 /* initialize the client library */
8 PMIx_Init(&myproc, NULL, 0);
9
10 /* get the #nodes in a different session */
11 sid = 12345;
12 PMIx_Info_Load(&info[0], PMIX_SESSION_INFO, NULL, PMIX_BOOL);
13 PMIx_Info_Load(&info[1], PMIX_SESSION_ID, &sid, PMIX_UINT32);
14 rc = PMIx_Get(NULL, PMIX_NUM_NODES, info, 2, &value);
```

C

15 6.2.1.2 Job-level information

16 Information regarding a job can be obtained by the methods detailed in Section 6.1.2. Example
17 requests are shown below:

C

```
18 pmix_info_t info;
19 pmix_value_t *value;
20 pmix_status_t rc;
21 pmix_proc_t myproc, wildcard;
22
23 /* initialize the client library */
24 PMIx_Init(&myproc, NULL, 0);
25
26 /* get the #apps in our job */
27 PMIX_PROC_LOAD(&wildcard, myproc.nspace, PMIX_RANK_WILDCARD);
28 rc = PMIx_Get(&wildcard, PMIX_JOB_NUM_APPS, NULL, 0, &value);
29
30 /* get the #nodes in our job */
31 PMIx_Info_Load(&info, PMIX_JOB_INFO, NULL, PMIX_BOOL);
32 rc = PMIx_Get(&wildcard, PMIX_NUM_NODES, &info, 1, &value);
```

C

1 6.2.1.3 Application-level information

2 Information regarding an application can be obtained by the methods described in Section 6.1.3.
3 Example requests are shown below:

```
4 pmix_info_t info;  
5 pmix_value_t *value;  
6 pmix_status_t rc;  
7 pmix_proc_t myproc, otherproc;  
8 uint32_t appsize, appnum;  
9  
10 /* initialize the client library */  
11 PMIx_Init(&myproc, NULL, 0);  
12  
13 /* get the #processes in our application */  
14 rc = PMIx_Get(&myproc, PMIX_APP_SIZE, NULL, 0, &value);  
15 appsize = value->data.uint32;  
16  
17 /* get the #nodes in an application containing "otherproc".  
18  * For this use-case, assume that we are in the first application  
19  * and we want the #nodes in the second application - use the  
20  * rank of the first process in that application, remembering  
21  * that ranks start at zero */  
22 PMIX_PROC_LOAD(&otherproc, myproc.nspace, appsize);  
23  
24 /* Since "otherproc" refers to a process in the second application,  
25  * we can simply mark that we want the info for this key from the  
26  * application realm */  
27 PMIx_Info_load(&info, PMIX_APP_INFO, NULL, PMIX_BOOL);  
28 rc = PMIx_Get(&otherproc, PMIX_NUM_NODES, &info, 1, &value);  
29  
30 /* alternatively, we can directly ask for the #nodes in  
31  * the second application in our job, again remembering that  
32  * application numbers start with zero. Since we are asking  
33  * for application realm information about a specific appnum  
34  * within our own namespace, the process identifier can be NULL */  
35 appnum = 1;  
36 PMIx_Info_load(&appinfo[0], PMIX_APP_INFO, NULL, PMIX_BOOL);  
37 PMIx_Info_load(&appinfo[1], PMIX_APPNUM, &appnum, PMIX_UINT32);  
38 rc = PMIx_Get(NULL, PMIX_NUM_NODES, appinfo, 2, &value);
```

1 **6.2.1.4 Process-level information**

2 Process-level information is accessed by providing the namespace and rank of the target process. In
3 the absence of any directive as to the level of information being requested, the PMIx library will
4 always return the process-level value. See Section 6.1.4 for details.

5 **6.2.1.5 Node-level information**

6 Information regarding a node within the system can be obtained by the methods described in
7 Section 6.1.5. Example requests are shown below:

```
8 pmix_info_t info[2];  
9 pmix_value_t *value;  
10 pmix_status_t rc;  
11 pmix_proc_t myproc, otherproc;  
12 uint32_t nodeid;  
  
13  
14 /* initialize the client library */  
15 PMIx_Init(&myproc, NULL, 0);  
16  
17 /* get the #procs on our node */  
18 rc = PMIx_Get(&myproc, PMIX_NODE_SIZE, NULL, 0, &value);  
19  
20 /* get the #slots on another node */  
21 PMIx_Info_load(&info[0], PMIX_NODE_INFO, NULL, PMIX_BOOL);  
22 PMIx_Info_load(&info[1], PMIX_HOSTNAME, "remotehost", PMIX_STRING);  
23 rc = PMIx_Get(NULL, PMIX_MAX_PROCS, info, 2, &value);  
24  
25 /* get the total #procs on the remote node - note that we don't  
26 * actually need to include the "PMIX_NODE_INFO" attribute here,  
27 * but (a) it does no harm and (b) it allowed us to simply reuse  
28 * the prior info array  
29 rc = PMIx_Get(NULL, PMIX_NODE_SIZE, info, 2, &value);
```


CHAPTER 7

Query Operations

1 This chapter presents mechanisms for generalized queries that access information about the host
2 environment and the system in general. The chapter presents the concept of a query followed by a
3 detailed explanation of the query APIs provided. The chapter compares the use of these APIs with
4 [PMIx_Get](#). The chapter concludes with detailed information about how to use the query interface
5 to access information about what PMIx APIs an implementation supports as well as what attributes
6 each supported API supports.

7 7.1 PMIx_Query_info

8 As the level of interaction between applications and the host SMS grows, so too does the need for
9 the application to query the SMS regarding its capabilities and state information. PMIx provides a
10 generalized query interface for this purpose, along with a set of standardized attribute keys to
11 support a range of requests. This includes requests to determine the status of scheduling queues and
12 active allocations, the scope of API and attribute support offered by the SMS, namespaces of active
13 jobs, location and information about a job's processes, and information regarding available
14 resources.

15 An example use-case for the [PMIx_Query_info_nb](#) API is to ensure clean job completion.
16 Time-shared systems frequently impose maximum run times when assigning jobs to resource
17 allocations. To shut down gracefully (e.g., to write a checkpoint before termination) it is necessary
18 for an application to periodically query the resource manager for the time remaining in its
19 allocation. This is especially true on systems for which allocation times may be shortened or
20 lengthened from the original time limit. Many resource managers provide APIs to dynamically
21 obtain this information, but each API is specific to the resource manager. PMIx supports this
22 use-case by defining an attribute key ([PMIX_TIME_REMAINING](#)) that can be used with the
23 [PMIx_Query_info_nb](#) interface to obtain the number of seconds remaining in the current job
24 allocation.

25 PMIx sometimes provides multiple methods by which an application can obtain information or
26 services. For this example, note that one could alternatively use the
27 [PMIx_Register_event_handler](#) API to register for an event indicating incipient job
28 termination, and then use the [PMIx_Job_control_nb](#) API to request that the host SMS
29 generate an event a specified amount of time prior to reaching the maximum run time.

1 7.1.1 Query Structure

2 A PMIx query structure is composed of one or more keys and a list of qualifiers which provide
3 additional information to describe the query. Keys which use the same qualifiers can be placed in
4 the same query for compactness, though it is permissible to put each key in its own query.

5 The `pmix_query_t` structure is used by the `PMIx_Query_info` APIs to describe a single
6 query operation.

PMIx v2.0

```
▼ _____ C _____ ▼  
7 typedef struct pmix_query {  
8     char **keys;  
9     pmix_info_t *qualifiers;  
10    size_t nqual;  
11 } pmix_query_t;  
▲ _____ C _____ ▲
```

12 where:

- 13 • *keys* is a **NULL**-terminated argv-style array of strings
- 14 • *qualifiers* is an array of `pmix_info_t` describing constraints on the query
- 15 • *nqual* is the number of elements in the *qualifiers* array

16 The following APIs support query of various session and environment values.

17 7.1.2 PMIx_Query_info

18 Summary

19 Query information about the system in general.

20 Format

PMIx v4.0

```
▼ _____ C _____ ▼  
21 pmix_status_t  
22 PMIx_Query_info(pmix_query_t queries[], size_t nqueries,  
23                 pmix_info_t *info[], size_t *ninfo);  
▲ _____ C _____ ▲
```

24 **IN** *queries*

25 Array of query structures (array of handles)

26 **IN** *nqueries*

27 Number of elements in the *queries* array (integer)

28 **INOUT** *info*

29 Address where a pointer to an array of `pmix_info_t` containing the results of the query can
30 be returned (memory reference)

INOUT *ninfo*

Address where the number of elements in *info* can be returned (handle)

A successful return indicates that all data was found and has been returned.

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- `PMIX_ERR_NOT_FOUND` None of the requested data was available.
- `PMIX_ERR_PARTIAL_SUCCESS` Some of the requested data was found. The *info* array shall contain an element for each query key that returned a value.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

If a value other than `PMIX_SUCCESS` or `PMIX_ERR_PARTIAL_SUCCESS` is returned, the *info* array shall be `NULL` and *ninfo* zero.

Required Attributes

A call to this API can specify multiple queries. Each query is composed of a list of keys and a list of attributes which can influence that query. PMIx libraries and host environments that support this API are required to support the following attributes which are specified on a per-query basis:

`PMIX_QUERY_REFRESH_CACHE` "`pmix.qry.rfsh`" (bool)

Retrieve updated information from server. NO QUALIFIERS.

`PMIX_SESSION_INFO` "`pmix.ssn.info`" (bool)

Return information regarding the session realm of the target process.

`PMIX_JOB_INFO` "`pmix.job.info`" (bool)

Return information regarding the job realm corresponding to the namespace in the target process' identifier.

`PMIX_APP_INFO` "`pmix.app.info`" (bool)

Return information regarding the application realm to which the target process belongs - the namespace of the target process serves to identify the job containing the target application. If information about an application other than the one containing the target process is desired, then the attribute array must contain a `PMIX_APPNUM` attribute identifying the desired target application. This is useful in cases where there are multiple applications and the mapping of processes to applications is unclear.

`PMIX_NODE_INFO` "`pmix.node.info`" (bool)

Return information from the node realm regarding the node upon which the specified process is executing. If information about a node other than the one containing the specified process is desired, then the attribute array must also contain either the `PMIX_NODEID` or `PMIX_HOSTNAME` attribute identifying the desired target. This is useful for requesting information about a specific node even if the identity of processes running on that node are not known.

1 **PMIX_PROC_INFO** "pmix.proc.info" (bool)
2 Return information regarding the target process. This attribute is technically not required as
3 the **PMIx_Get** API specifically identifies the target process in its parameters. However, it is
4 included here for completeness.

5 **PMIX_PROCID** "pmix.procid" (pmix_proc_t)
6 The caller's process identifier. The value returned is identical to what **PMIx_Init** or
7 **PMIx_tool_init** provides. The process identifier in the **PMIx_Get** call is ignored when
8 requesting this key. In this context, specifies the process ID whose information is being
9 requested - e.g., a query asking for the **pmix_proc_info_t** of a specified process. Only
10 required when the request is for information on a specific process.

11 **PMIX_NAMESPACE** "pmix.namespace" (char*)
12 Namespace of the job - may be a numerical value expressed as a string, but is often an
13 alphanumeric string carrying information solely of use to the system. Required to be unique
14 within the scope of the host environment. One cannot retrieve the namespace of an arbitrary
15 process since that would require already knowing the namespace of that process. However, a
16 process' own namespace can be retrieved by passing a NULL value of *proc* to **PMIx_Get**.
17 Specifies the namespace of the process whose information is being requested. Must be
18 accompanied by the **PMIX_RANK** attribute. Only required when the request is for
19 information on a specific process.

20 **PMIX_RANK** "pmix.rank" (pmix_rank_t)
21 Process rank within the job, starting from zero. Specifies the rank of the process whose
22 information is being requested. Must be accompanied by the **PMIX_NAMESPACE** attribute.
23 Only required when the request is for information on a specific process.

24 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "pmix.qry.attrs" (bool)
25 Query list of supported attributes for specified APIs. REQUIRED QUALIFIERS: one or
26 more of **PMIX_CLIENT_FUNCTIONS**, **PMIX_SERVER_FUNCTIONS**,
27 **PMIX_TOOL_FUNCTIONS**, and **PMIX_HOST_FUNCTIONS**.

28 **PMIX_CLIENT_ATTRIBUTES** "pmix.client.attrs" (bool)
29 Request attributes supported by the PMIx client library.

30 **PMIX_SERVER_ATTRIBUTES** "pmix.srvr.attrs" (bool)
31 Request attributes supported by the PMIx server library.

32 **PMIX_HOST_ATTRIBUTES** "pmix.host.attrs" (bool)
33 Request attributes supported by the host environment.

34 **PMIX_TOOL_ATTRIBUTES** "pmix.setup.env" (bool)
35 Request attributes supported by the PMIx tool library functions.

36 Note that inclusion of both the **PMIX_PROCID** directive and either the **PMIX_NAMESPACE** or the
37 **PMIX_RANK** attribute will return a **PMIX_ERR_BAD_PARAM** result, and that the inclusion of a
38 process identifier must apply to all keys in that **pmix_query_t**. Queries for information on

1 multiple specific processes therefore requires submitting multiple `pmix_query_t` structures,
2 each referencing one process. Directives which are not applicable to a key are ignored.

3 An implementation is not required to support any particular keys. If a key is unsupported, the
4 implementation should handle that key in the same way that it is required to handle a key which it
5 cannot find. The following keys may be specified in a query:

6 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "`pmix.qry.attrs`" (`bool`)

7 Query list of supported attributes for specified APIs. REQUIRED QUALIFIERS: one or
8 more of **PMIX_CLIENT_FUNCTIONS**, **PMIX_SERVER_FUNCTIONS**,
9 **PMIX_TOOL_FUNCTIONS**, and **PMIX_HOST_FUNCTIONS**.

10 **PMIX_QUERY_NAMESPACES** "`pmix.qry.ns`" (`char*`)

11 Request a comma-delimited list of active namespaces. NO QUALIFIERS.

12 **PMIX_QUERY_JOB_STATUS** "`pmix.qry.jst`" (`pmix_status_t`)

13 Status of a specified, currently executing job. REQUIRED QUALIFIER: **PMIX_NAMESPACE**
14 indicating the namespace whose status is being queried.

15 **PMIX_QUERY_QUEUE_LIST** "`pmix.qry.qlst`" (`char*`)

16 Request a comma-delimited list of scheduler queues. NO QUALIFIERS.

17 **PMIX_QUERY_QUEUE_STATUS** "`pmix.qry.qst`" (`char*`)

18 Returns status of a specified scheduler queue, expressed as a string. OPTIONAL
19 QUALIFIERS: **PMIX_ALLOC_QUEUE** naming specific queue whose status is being
20 requested.

21 **PMIX_QUERY_PROC_TABLE** "`pmix.qry.phtable`" (`char*`)

22 Returns a (`pmix_data_array_t`) array of `pmix_proc_info_t`, one entry for each
23 process in the specified namespace, ordered by process job rank. REQUIRED QUALIFIER:
24 **PMIX_NAMESPACE** indicating the namespace whose process table is being queried.

25 **PMIX_QUERY_LOCAL_PROC_TABLE** "`pmix.qry.lhtable`" (`char*`)

26 Returns a (`pmix_data_array_t`) array of `pmix_proc_info_t`, one entry for each
27 process in the specified namespace executing on the same node as the requester, ordered by
28 process job rank. REQUIRED QUALIFIER: **PMIX_NAMESPACE** indicating the namespace
29 whose local process table is being queried. OPTIONAL QUALIFIER: **PMIX_HOSTNAME**
30 indicating the host whose local process table is being queried. By default, the query assumes
31 that the host upon which the request was made is to be used.

32 **PMIX_QUERY_SPAWN_SUPPORT** "`pmix.qry.spawn`" (`bool`)

33 Return a comma-delimited list of supported spawn attributes. NO QUALIFIERS.

34 **PMIX_QUERY_DEBUG_SUPPORT** "`pmix.qry.debug`" (`bool`)

35 Return a comma-delimited list of supported debug attributes. NO QUALIFIERS.

36 **PMIX_QUERY_MEMORY_USAGE** "`pmix.qry.mem`" (`bool`)

1 Return information on memory usage for the processes indicated in the qualifiers.
2 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** and **PMIX_RANK**, or **PMIX_PROCID** of
3 specific process(es) whose memory usage is being requested.

4 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)
5 Report only average values for sampled information. NO QUALIFIERS.

6 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)
7 Report minimum and maximum values. NO QUALIFIERS.

8 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
9 String identifier of the allocation whose status is being requested. NO QUALIFIERS.

10 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)
11 Query number of seconds (**uint32_t**) remaining in allocation for the specified namespace.
12 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** of the namespace whose info is being
13 requested (defaults to allocation containing the caller).

14 **PMIX_SERVER_URI** "pmix.srvr.uri" (char*)
15 URI of the PMIx server to be contacted. Requests the URI of the specified PMIx server's
16 PMIx connection. Defaults to requesting the information for the local PMIx server.

17 **PMIX_CLIENT_AVG_MEMORY** "pmix.cl.mem.avg" (float)
18 Average Megabytes of memory used by client processes on node. OPTIONAL
19 QUALIFIERS: **PMIX_HOSTNAME** or **PMIX_NODEID** (defaults to caller's node).

20 **PMIX_DAEMON_MEMORY** "pmix.dmn.mem" (float)
21 Megabytes of memory currently used by the RM daemon on the node. OPTIONAL
22 QUALIFIERS: **PMIX_HOSTNAME** or **PMIX_NODEID** (defaults to caller's node).

23 **PMIX_QUERY_AUTHORIZATIONS** "pmix.qry.auths" (bool)
24 Return operations the PMIx tool is authorized to perform. NO QUALIFIERS.

25 **PMIX_PROC_PID** "pmix.ppid" (pid_t)
26 Operating system PID of specified process.

27 **PMIX_PROC_STATE_STATUS** "pmix.proc.state" (pmix_proc_state_t)
28 State of the specified process as of the last report - may not be the actual current state based
29 on update rate.

30 Description

31 Query information about the system in general. This can include a list of active namespaces, fabric
32 topology, etc. Also can be used to query node-specific info such as the list of peers executing on a
33 given node. The host environment is responsible for exercising appropriate access control on the
34 information.

35 The returned *status* indicates if requested data was found or not. The returned *info* array will
36 contain a **PMIX_QUERY_RESULTS** element for each query of the *queries* array. If qualifiers were
37 included in the query, then the first element of each results array shall contain the

1 **PMIX_QUERY_QUALIFIERS** key with a **pmix_data_array_t** containing the qualifiers. The
2 remaining **pmix_info_t** shall contain the results of the query, one entry for each key that was
3 found. Note that duplicate keys in the *queries* array shall result in duplicate responses within the
4 constraints of the accompanying qualifiers. The caller is responsible for releasing the returned array.

Advice to PMIx library implementers

5 It is recommended that information returned from **PMIx_Query_info** be locally cached so that
6 retrieval by subsequent calls to **PMIx_Get**, **PMIx_Query_info**, or **PMIx_Query_info_nb**
7 can succeed with minimal overhead. The local cache shall be checked prior to querying the PMIx
8 server and/or the host environment. Queries that include the **PMIX_QUERY_REFRESH_CACHE**
9 attribute shall bypass the local cache and retrieve a new value for the query, refreshing the values in
10 the cache upon return.

7.1.3 **PMIx_Query_info_nb**

Summary

13 Query information about the system in general.

Format

14 *PMIx v2.0*

```
15 pmix_status_t  
16 PMIx_Query_info_nb(pmix_query_t queries[], size_t nqueries,  
17 pmix_info_cbfunc_t cbfunc, void *cbdata);
```

IN **queries**

19 Array of query structures (array of handles)

IN **nqueries**

21 Number of elements in the *queries* array (integer)

IN **cbfunc**

23 Callback function **pmix_info_cbfunc_t** (function reference)

IN **cbdata**

25 Data to be passed to the callback function (memory reference)

26 A successful return indicates that the request has been accepted for processing. The provided
27 callback function will only be executed upon successful return of the operation. Note that the
28 library must not invoke the callback function prior to returning from the API.

29 Returns **PMIX_SUCCESS** or a negative value indicating the error.

30 If executed, the status returned in the provided callback function will be one of the following
31 constants:

- **PMIX_SUCCESS** All data was found and has been returned.

- 1 • **PMIX_ERR_NOT_FOUND** None of the requested data was available. The *info* array will be
2 **NULL** and *ninfo* zero.
- 3 • **PMIX_ERR_PARTIAL_SUCCESS** Some of the requested data was found. The *info* array shall
4 contain an element for each query key that returned a value.
- 5 • **PMIX_ERR_NOT_SUPPORTED** The host RM does not support this function. The *info* array will
6 be **NULL** and *ninfo* zero.
- 7 • a non-zero PMIx error constant indicating a reason for the request's failure. The *info* array will
8 be **NULL** and *ninfo* zero.

Required Attributes

9 PMIx libraries and host environments that support this API are required to support the following
10 attributes:

11 **PMIX_QUERY_REFRESH_CACHE** "pmix.qry.rfsh" (bool)

12 Retrieve updated information from server. NO QUALIFIERS.

13 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)

14 Return information regarding the session realm of the target process.

15 **PMIX_JOB_INFO** "pmix.job.info" (bool)

16 Return information regarding the job realm corresponding to the namespace in the target
17 process' identifier.

18 **PMIX_APP_INFO** "pmix.app.info" (bool)

19 Return information regarding the application realm to which the target process belongs - the
20 namespace of the target process serves to identify the job containing the target application. If
21 information about an application other than the one containing the target process is desired,
22 then the attribute array must contain a **PMIX_APPNUM** attribute identifying the desired
23 target application. This is useful in cases where there are multiple applications and the
24 mapping of processes to applications is unclear.

25 **PMIX_NODE_INFO** "pmix.node.info" (bool)

26 Return information from the node realm regarding the node upon which the specified
27 process is executing. If information about a node other than the one containing the specified
28 process is desired, then the attribute array must also contain either the **PMIX_NODEID** or
29 **PMIX_HOSTNAME** attribute identifying the desired target. This is useful for requesting
30 information about a specific node even if the identity of processes running on that node are
31 not known.

32 **PMIX_PROC_INFO** "pmix.proc.info" (bool)

33 Return information regarding the target process. This attribute is technically not required as
34 the **PMIx_Get** API specifically identifies the target process in its parameters. However, it is
35 included here for completeness.

36 **PMIX_PROCID** "pmix.procid" (pmix_proc_t)

1 The caller's process identifier. The value returned is identical to what `PMIx_Init` or
2 `PMIx_tool_init` provides. The process identifier in the `PMIx_Get` call is ignored when
3 requesting this key. In this context, specifies the process ID whose information is being
4 requested - e.g., a query asking for the `pmix_proc_info_t` of a specified process. Only
5 required when the request is for information on a specific process.

6 **PMIX_NAMESPACE** "`pmix.namespace`" (`char*`)

7 Namespace of the job - may be a numerical value expressed as a string, but is often an
8 alphanumeric string carrying information solely of use to the system. Required to be unique
9 within the scope of the host environment. One cannot retrieve the namespace of an arbitrary
10 process since that would require already knowing the namespace of that process. However, a
11 process' own namespace can be retrieved by passing a NULL value of `proc` to `PMIx_Get`.
12 Specifies the namespace of the process whose information is being requested. Must be
13 accompanied by the **PMIX_RANK** attribute. Only required when the request is for
14 information on a specific process.

15 **PMIX_RANK** "`pmix.rank`" (`pmix_rank_t`)

16 Process rank within the job, starting from zero. Specifies the rank of the process whose
17 information is being requested. Must be accompanied by the **PMIX_NAMESPACE** attribute.
18 Only required when the request is for information on a specific process.

19 **PMIX_CLIENT_ATTRIBUTES** "`pmix.client.attrs`" (`bool`)

20 Request attributes supported by the PMIx client library.

21 **PMIX_SERVER_ATTRIBUTES** "`pmix.srvr.attrs`" (`bool`)

22 Request attributes supported by the PMIx server library.

23 **PMIX_HOST_ATTRIBUTES** "`pmix.host.attrs`" (`bool`)

24 Request attributes supported by the host environment.

25 **PMIX_TOOL_ATTRIBUTES** "`pmix.setup.env`" (`bool`)

26 Request attributes supported by the PMIx tool library functions.

27 Note that inclusion of both the **PMIX_PROCID** directive and either the **PMIX_NAMESPACE** or the
28 **PMIX_RANK** attribute will return a **PMIX_ERR_BAD_PARAM** result, and that the inclusion of a
29 process identifier must apply to all keys in that `pmix_query_t`. Queries for information on
30 multiple specific processes therefore requires submitting multiple `pmix_query_t` structures,
31 each referencing one process. Directives which are not applicable to a key are ignored.

▲-----▲
32 An implementation is not required to support any particular keys. If a key is unsupported, the
33 implementation should handle that key in the same way that it is required to handle a key which it
34 cannot find. The following keys may be specified in a query:

35 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "`pmix.qry.attrs`" (`bool`)

1 Query list of supported attributes for specified APIs. REQUIRED QUALIFIERS: one or
2 more of **PMIX_CLIENT_FUNCTIONS**, **PMIX_SERVER_FUNCTIONS**,
3 **PMIX_TOOL_FUNCTIONS**, and **PMIX_HOST_FUNCTIONS**.

4 **PMIX_QUERY_NAMESPACES** "pmix.qry.ns" (char*)

5 Request a comma-delimited list of active namespaces. NO QUALIFIERS.

6 **PMIX_QUERY_JOB_STATUS** "pmix.qry.jst" (pmix_status_t)

7 Status of a specified, currently executing job. REQUIRED QUALIFIER: **PMIX_NAMESPACE**
8 indicating the namespace whose status is being queried.

9 **PMIX_QUERY_QUEUE_LIST** "pmix.qry.qlst" (char*)

10 Request a comma-delimited list of scheduler queues. NO QUALIFIERS.

11 **PMIX_QUERY_QUEUE_STATUS** "pmix.qry.qst" (char*)

12 Returns status of a specified scheduler queue, expressed as a string. OPTIONAL
13 QUALIFIERS: **PMIX_ALLOC_QUEUE** naming specific queue whose status is being
14 requested.

15 **PMIX_QUERY_PROC_TABLE** "pmix.qry.phtable" (char*)

16 Returns a (**pmix_data_array_t**) array of **pmix_proc_info_t**, one entry for each
17 process in the specified namespace, ordered by process job rank. REQUIRED QUALIFIER:
18 **PMIX_NAMESPACE** indicating the namespace whose process table is being queried.

19 **PMIX_QUERY_LOCAL_PROC_TABLE** "pmix.qry.lhtable" (char*)

20 Returns a (**pmix_data_array_t**) array of **pmix_proc_info_t**, one entry for each
21 process in the specified namespace executing on the same node as the requester, ordered by
22 process job rank. REQUIRED QUALIFIER: **PMIX_NAMESPACE** indicating the namespace
23 whose local process table is being queried. OPTIONAL QUALIFIER: **PMIX_HOSTNAME**
24 indicating the host whose local process table is being queried. By default, the query assumes
25 that the host upon which the request was made is to be used.

26 **PMIX_QUERY_SPAWN_SUPPORT** "pmix.qry.spawn" (bool)

27 Return a comma-delimited list of supported spawn attributes. NO QUALIFIERS.

28 **PMIX_QUERY_DEBUG_SUPPORT** "pmix.qry.debug" (bool)

29 Return a comma-delimited list of supported debug attributes. NO QUALIFIERS.

30 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)

31 Return information on memory usage for the processes indicated in the qualifiers.
32 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** and **PMIX_RANK**, or **PMIX_PROCID** of
33 specific process(es) whose memory usage is being requested.

34 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)

35 Report only average values for sampled information. NO QUALIFIERS.

36 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)

37 Report minimum and maximum values. NO QUALIFIERS.

1 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
 2 String identifier of the allocation whose status is being requested. NO QUALIFIERS.

3 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)
 4 Query number of seconds (uint32_t) remaining in allocation for the specified namespace.
 5 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** of the namespace whose info is being
 6 requested (defaults to allocation containing the caller).

7 **PMIX_SERVER_URI** "pmix.srvr.uri" (char*)
 8 URI of the PMIx server to be contacted. Requests the URI of the specified PMIx server's
 9 PMIx connection. Defaults to requesting the information for the local PMIx server.

10 **PMIX_CLIENT_AVG_MEMORY** "pmix.cl.mem.avg" (float)
 11 Average Megabytes of memory used by client processes on node. OPTIONAL
 12 QUALIFIERS: **PMIX_HOSTNAME** or **PMIX_NODEID** (defaults to caller's node).

13 **PMIX_DAEMON_MEMORY** "pmix.dmn.mem" (float)
 14 Megabytes of memory currently used by the RM daemon on the node. OPTIONAL
 15 QUALIFIERS: **PMIX_HOSTNAME** or **PMIX_NODEID** (defaults to caller's node).

16 **PMIX_QUERY_AUTHORIZATIONS** "pmix.qry.auths" (bool)
 17 Return operations the PMIx tool is authorized to perform. NO QUALIFIERS.

18 **PMIX_PROC_PID** "pmix.ppid" (pid_t)
 19 Operating system PID of specified process.

20 **PMIX_PROC_STATE_STATUS** "pmix.proc.state" (pmix_proc_state_t)
 21 State of the specified process as of the last report - may not be the actual current state based
 22 on update rate.

23 Description

24 Non-blocking form of the [PMIx_Query_info](#) API.

25 7.1.4 Query keys

26 The following keys may be queried using the [PMIx_Query_info](#) and
 27 [PMIx_Query_info_nb](#) APIs:

28 **PMIX_QUERY_SUPPORTED_KEYS** "pmix.qry.keys" (char*)
 29 Returns comma-delimited list of keys supported by the query function. NO QUALIFIERS.

30 **PMIX_QUERY_SUPPORTED_QUALIFIERS** "pmix.qry.qual" (char*)
 31 Return comma-delimited list of qualifiers supported by a query on the provided key, instead
 32 of actually performing the query on the key. NO QUALIFIERS.

33 **PMIX_QUERY_NAMESPACES** "pmix.qry.ns" (char*)
 34 Request a comma-delimited list of active namespaces. NO QUALIFIERS.

35 **PMIX_QUERY_NAMESPACE_INFO** "pmix.qry.nsinfo" (pmix_data_array_t*)

1 Return an array of active namespace information - each element will itself contain an array
2 including the namespace plus the command line of the application executing within it.
3 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** of specific namespace whose info is being
4 requested.

5 **PMIX_QUERY_JOB_STATUS** "pmix.qry.jst" (pmix_status_t)

6 Status of a specified, currently executing job. REQUIRED QUALIFIER: **PMIX_NAMESPACE**
7 indicating the namespace whose status is being queried.

8 **PMIX_QUERY_QUEUE_LIST** "pmix.qry qlst" (char*)

9 Request a comma-delimited list of scheduler queues. NO QUALIFIERS.

10 **PMIX_QUERY_QUEUE_STATUS** "pmix.qry.qst" (char*)

11 Returns status of a specified scheduler queue, expressed as a string. OPTIONAL
12 QUALIFIERS: **PMIX_ALLOC_QUEUE** naming specific queue whose status is being
13 requested.

14 **PMIX_QUERY_PROC_TABLE** "pmix.qry.phtable" (char*)

15 Returns a (pmix_data_array_t) array of pmix_proc_info_t, one entry for each
16 process in the specified namespace, ordered by process job rank. REQUIRED QUALIFIER:
17 **PMIX_NAMESPACE** indicating the namespace whose process table is being queried.

18 **PMIX_QUERY_LOCAL_PROC_TABLE** "pmix.qry.lhtable" (char*)

19 Returns a (pmix_data_array_t) array of pmix_proc_info_t, one entry for each
20 process in the specified namespace executing on the same node as the requester, ordered by
21 process job rank. REQUIRED QUALIFIER: **PMIX_NAMESPACE** indicating the namespace
22 whose local process table is being queried. OPTIONAL QUALIFIER: **PMIX_HOSTNAME**
23 indicating the host whose local process table is being queried. By default, the query assumes
24 that the host upon which the request was made is to be used.

25 **PMIX_QUERY_AUTHORIZATIONS** "pmix.qry.auths" (bool)

26 Return operations the PMIX tool is authorized to perform. NO QUALIFIERS.

27 **PMIX_QUERY_SPAWN_SUPPORT** "pmix.qry.spawn" (bool)

28 Return a comma-delimited list of supported spawn attributes. NO QUALIFIERS.

29 **PMIX_QUERY_DEBUG_SUPPORT** "pmix.qry.debug" (bool)

30 Return a comma-delimited list of supported debug attributes. NO QUALIFIERS.

31 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)

32 Return information on memory usage for the processes indicated in the qualifiers.
33 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** and **PMIX_RANK**, or **PMIX_PROCID** of
34 specific process(es) whose memory usage is being requested.

35 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)

36 Query number of seconds (uint32_t) remaining in allocation for the specified namespace.
37 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** of the namespace whose info is being
38 requested (defaults to allocation containing the caller).

39 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "pmix.qry.attrs" (bool)

40 Query list of supported attributes for specified APIs. REQUIRED QUALIFIERS: one or
41 more of **PMIX_CLIENT_FUNCTIONS**, **PMIX_SERVER_FUNCTIONS**,
42 **PMIX_TOOL_FUNCTIONS**, and **PMIX_HOST_FUNCTIONS**.

1 **PMIX_QUERY_NUM_PSETS** "pmix.qry.psetnum" (`size_t`)
 2 Return the number of process sets defined in the specified range (defaults to
 3 **PMIX_RANGE_SESSION**).

4 **PMIX_QUERY_PSET_NAMES** "pmix.qry.psets" (`pmix_data_array_t*`)
 5 Return a `pmix_data_array_t` containing an array of strings of the process set names
 6 defined in the specified range (defaults to **PMIX_RANGE_SESSION**).

7 **PMIX_QUERY_PSET_MEMBERSHIP** "pmix.qry.pmems" (`pmix_data_array_t*`)
 8 Return an array of `pmix_proc_t` containing the members of the specified process set.

9 **PMIX_QUERY_AVAIL_SERVERS** "pmix.qry.asrvrs" (`pmix_data_array_t*`)
 10 Return an array of `pmix_info_t`, each element itself containing a
 11 **PMIX_SERVER_INFO_ARRAY** entry holding all available data for a server on this node to
 12 which the caller might be able to connect.

13 These keys are used to query memory available and used in the system.

14 **PMIX_AVAIL_PHYS_MEMORY** "pmix.pmem" (`uint64_t`)
 15 Total available physical memory on a node. OPTIONAL QUALIFERS: **PMIX_HOSTNAME**
 16 or **PMIX_NODEID** (defaults to caller's node).

17 **PMIX_DAEMON_MEMORY** "pmix.dmn.mem" (`float`)
 18 Megabytes of memory currently used by the RM daemon on the node. OPTIONAL
 19 QUALIFERS: **PMIX_HOSTNAME** or **PMIX_NODEID** (defaults to caller's node).

20 **PMIX_CLIENT_AVG_MEMORY** "pmix.cl.mem.avg" (`float`)
 21 Average Megabytes of memory used by client processes on node. OPTIONAL
 22 QUALIFERS: **PMIX_HOSTNAME** or **PMIX_NODEID** (defaults to caller's node).

23 7.1.5 Query attributes

24 Attributes used to direct behavior of the **PMIx_Query_info** and **PMIx_Query_info_nb**
 25 APIs:

26 **PMIX_QUERY_RESULTS** "pmix.qry.res" (`pmix_data_array_t`)
 27 Contains an array of query results for a given `pmix_query_t` passed to the
 28 **PMIx_Query_info** APIs. If qualifiers were included in the query, then the first element
 29 of the array shall be the **PMIX_QUERY_QUALIFIERS** attribute containing those qualifiers.
 30 Each of the remaining elements of the array is a `pmix_info_t` containing the query key
 31 and the corresponding value returned by the query. This attribute is solely for reporting
 32 purposes and cannot be used in **PMIx_Get** or other query operations.

33 **PMIX_QUERY_QUALIFIERS** "pmix.qryquals" (`pmix_data_array_t`)
 34 Contains an array of qualifiers that were included in the query that produced the provided
 35 results. This attribute is solely for reporting purposes and cannot be used in **PMIx_Get** or
 36 other query operations.

37 **PMIX_QUERY_REFRESH_CACHE** "pmix.qry.rfsh" (`bool`)
 38 Retrieve updated information from server. NO QUALIFIERS.

1 **PMIX_QUERY_LOCAL_ONLY** "pmix.qry.local" (bool)
 2 Constrain the query to local information only. NO QUALIFIERS.
 3 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)
 4 Report only average values for sampled information. NO QUALIFIERS.
 5 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)
 6 Report minimum and maximum values. NO QUALIFIERS.
 7 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
 8 String identifier of the allocation whose status is being requested. NO QUALIFIERS.
 9 **PMIX_SERVER_INFO_ARRAY** "pmix.srv.arr" (pmix_data_array_t)
 10 Array of **pmix_info_t** about a given server, starting with its **PMIX_NAMESPACE** and
 11 including at least one of the rendezvous-required pieces of information.

12 The following attributes are used as qualifiers in queries regarding attribute support within the
 13 PMIx implementation and/or the host environment: **PMIX_CLIENT_FUNCTIONS**

14 **"pmix.client.fns"** (bool)
 15 Request a list of functions supported by the PMIx client library.
 16 **PMIX_CLIENT_ATTRIBUTES** "pmix.client.attrs" (bool)
 17 Request attributes supported by the PMIx client library.
 18 **PMIX_SERVER_FUNCTIONS** "pmix.srvr.fns" (bool)
 19 Request a list of functions supported by the PMIx server library.
 20 **PMIX_SERVER_ATTRIBUTES** "pmix.srvr.attrs" (bool)
 21 Request attributes supported by the PMIx server library.
 22 **PMIX_HOST_FUNCTIONS** "pmix.srvr.fns" (bool)
 23 Request a list of functions supported by the host environment.
 24 **PMIX_HOST_ATTRIBUTES** "pmix.host.attrs" (bool)
 25 Request attributes supported by the host environment.
 26 **PMIX_TOOL_FUNCTIONS** "pmix.tool.fns" (bool)
 27 Request a list of functions supported by the PMIx tool library.
 28 **PMIX_TOOL_ATTRIBUTES** "pmix.setup.env" (bool)
 29 Request attributes supported by the PMIx tool library functions.

30 7.1.5.1 Query structure support macros

31 The following macros are provided to support the **pmix_query_t** structure.

32 **Static initializer for the query structure**

33 **Provisional** Provide a static initializer for the **pmix_query_t** fields.

34 *PMIx v4.2* 

PMIX_QUERY_STATIC_INIT 

1 **Initialize the query structure**

2 Initialize the `pmix_query_t` fields



3 **PMIX_QUERY_CONSTRUCT (m)**



4 **IN** m

5 Pointer to the structure to be initialized (pointer to `pmix_query_t`)

6 **Destruct the query structure**

7 Destruct the `pmix_query_t` fields

PMIx v2.0



8 **PMIX_QUERY_DESTRUCT (m)**



9 **IN** m

10 Pointer to the structure to be destructed (pointer to `pmix_query_t`)

11 **Create a query array**

12 Allocate and initialize an array of `pmix_query_t` structures

PMIx v2.0



13 **PMIX_QUERY_CREATE (m, n)**



14 **INOUT** m

15 Address where the pointer to the array of `pmix_query_t` structures shall be stored (handle)

16 **IN** n

17 Number of structures to be allocated (`size_t`)

18 **Free a query structure**

19 Release a `pmix_query_t` structure

PMIx v4.0




20 **PMIX_QUERY_RELEASE (m)**



21 **IN** m


22 Pointer to a `pmix_query_t` structure (handle)

1 **Free a query array**
2 Release an array of `pmix_query_t` structures



3 **PMIX_QUERY_FREE**(m, n)
4 **IN** m
5 Pointer to the array of `pmix_query_t` structures (handle)
6 **IN** n
7 Number of structures in the array (`size_t`)

8 **Create the info array of query qualifiers**
9 Create an array of `pmix_info_t` structures for passing query qualifiers, updating the `nqual` field
10 of the `pmix_query_t` structure.



11 **PMIX_QUERY_QUALIFIERS_CREATE**(m, n)
12 **IN** m
13 Pointer to the `pmix_query_t` structure (handle)
14 **IN** n
15 Number of qualifiers to be allocated (`size_t`)

16 7.2 PMIx_Resolve_peers

17 There are a number of common queries for which PMIx provides convenience routines. These APIs
18 provide simplified access to commonly requested queries. Due to their simplified interface, these
19 APIs cannot be customized through the use of attributes. If a more specialized version of these
20 queries are required, similar functionality can often be accessed through the `PMIx_Query_info`
21 or `PMIx_Query_info_nb` APIs.

22 Summary

23 Obtain the array of processes within the specified namespace that are executing on a given node.

Format

```
pmix_status_t
PMIx_Resolve_peers(const char *nodename,
const pmix_namespace_t nspace,
pmix_proc_t **procs, size_t *nprocs);
```

IN *nodename*

Name of the node to query - **NULL** can be used to denote the current local node (string)

IN *nspace*

namespace (string)

OUT *procs*

Array of process structures (array of handles)

OUT *nprocs*

Number of elements in the *procs* array (integer)

Returns **PMIX_SUCCESS** or a negative value indicating the error.

Description

Given a *nodename*, return the array of processes within the specified *nspace* that are executing on that node. If the *nspace* is **NULL**, then all processes on the node will be returned. If the specified node does not currently host any processes, then the returned array will be **NULL**, and *nprocs* will be zero. The caller is responsible for releasing the *procs* array when done with it. The **PMIX_PROC_FREE** macro is provided for this purpose.

7.2.1 PMIx_Resolve_nodes

Summary

Return a list of nodes hosting processes within the given namespace.

Format

```
pmix_status_t
PMIx_Resolve_nodes(const char *nspace, char **nodelist);
```

IN *nspace*

Namespace (string)

OUT *nodelist*

Comma-delimited list of nodenames (string)

Returns **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant.

Description

Given a *nospace*, return the list of nodes hosting processes within that namespace. The returned string will contain a comma-delimited list of nodenames. The caller is responsible for releasing the string when done with it.

7.3 Using Get vs Query

Both **PMIx_Get** and **PMIx_Query_info** can be used to retrieve information about the system. In general, the *get* operation should be used to retrieve:

- information provided by the host environment at time of job start. This includes information on the number of processes in the job, their location, and possibly their communication endpoints.
- information posted by processes via the **PMIx_Put** function.

This information is largely considered to be *static*, although this will not necessarily be true for environments supporting dynamic programming models or fault tolerance. Note that the **PMIx_Get** function only accesses information about execution environments - i.e., its scope is limited to values pertaining to a specific *session*, *job*, *application*, *process*, or *node*. It cannot be used to obtain information about areas such as the status of queues in the WLM.

In contrast, the *query* option should be used to access:

- system-level information (such as the available WLM queues) that would generally not be included in job-level information provided at job start.
- dynamic information such as application and queue status, and resource utilization statistics. Note that the **PMIX_QUERY_REFRESH_CACHE** attribute must be provided on each query to ensure current data is returned.
- information created post job start, such as process tables.
- information requiring more complex search criteria than supported by the simpler **PMIx_Get** API.
- queries focused on retrieving multi-attribute blocks of data with a single request, thus bypassing the single-key limitation of the **PMIx_Get** API.

In theory, all information can be accessed via **PMIx_Query_info** as the local cache is typically the same datastore searched by **PMIx_Get**. However, in practice, the overhead associated with the *query* operation may (depending upon implementation) be higher than the simpler *get* operation due to the need to construct and process the more complex **pmix_query_t** structure. Thus, requests for a single key value are likely to be accomplished faster with **PMIx_Get** versus the *query* operation.

1 7.4 Accessing attribute support information

2 Information as to which attributes are supported by either the PMIx implementation or its host
3 environment can be obtained via the [PMIx_Query_info](#) APIs. The
4 [PMIX_QUERY_ATTRIBUTE_SUPPORT](#) attribute must be listed as the first entry in the *keys* field
5 of the [pmix_query_t](#) structure, followed by the name of the function whose attribute support is
6 being requested - support for multiple functions can be requested simultaneously by simply adding
7 the function names to the array of *keys*. Function names *must* be given as user-level API names -
8 e.g., “PMIx_Get”, “PMIx_server_setup_application”, or “PMIx_tool_attach_to_server”.

9 The desired levels of attribute support are provided as qualifiers. Multiple levels can be requested
10 simultaneously by simply adding elements to the *qualifiers* array. Each qualifier should contain the
11 desired level attribute with the boolean value set to indicate whether or not that level is to be
12 included in the returned information. Failure to provide any levels is equivalent to a request for all
13 levels. Supported levels include:

- 14 • [PMIX_CLIENT_FUNCTIONS](#) "[pmix.client.fns](#)" ([bool](#))
15 Request a list of functions supported by the PMIx client library.
- 16 • [PMIX_CLIENT_ATTRIBUTES](#) "[pmix.client.attrs](#)" ([bool](#))
17 Request attributes supported by the PMIx client library.
- 18 • [PMIX_SERVER_FUNCTIONS](#) "[pmix.srvr.fns](#)" ([bool](#))
19 Request a list of functions supported by the PMIx server library.
- 20 • [PMIX_SERVER_ATTRIBUTES](#) "[pmix.srvr.attrs](#)" ([bool](#))
21 Request attributes supported by the PMIx server library.
- 22 • [PMIX_HOST_FUNCTIONS](#) "[pmix.srvr.fns](#)" ([bool](#))
23 Request a list of functions supported by the host environment.
- 24 • [PMIX_HOST_ATTRIBUTES](#) "[pmix.host.attrs](#)" ([bool](#))
25 Request attributes supported by the host environment.
- 26 • [PMIX_TOOL_FUNCTIONS](#) "[pmix.tool.fns](#)" ([bool](#))
27 Request a list of functions supported by the PMIx tool library.
- 28 • [PMIX_TOOL_ATTRIBUTES](#) "[pmix.setup.env](#)" ([bool](#))
29 Request attributes supported by the PMIx tool library functions.

30 Unlike other queries, queries for attribute support can result in the number of returned
31 [pmix_info_t](#) structures being different from the number of queries. Each element in the
32 returned array will correspond to a pair of specified attribute level and function in the query, where
33 the *key* is the function and the *value* contains a [pmix_data_array_t](#) of [pmix_info_t](#). Each
34 element of the array is marked by a *key* indicating the requested attribute *level* with a *value*
35 composed of a [pmix_data_array_t](#) of [pmix_regattr_t](#), each describing a supported
36 attribute for that function, as illustrated in Fig. 7.1 below where the requestor asked for supported

1
2

attributes of **PMIx_Get** at the *client* and *server* levels, plus attributes of **PMIx_Allocation_request** at all levels.

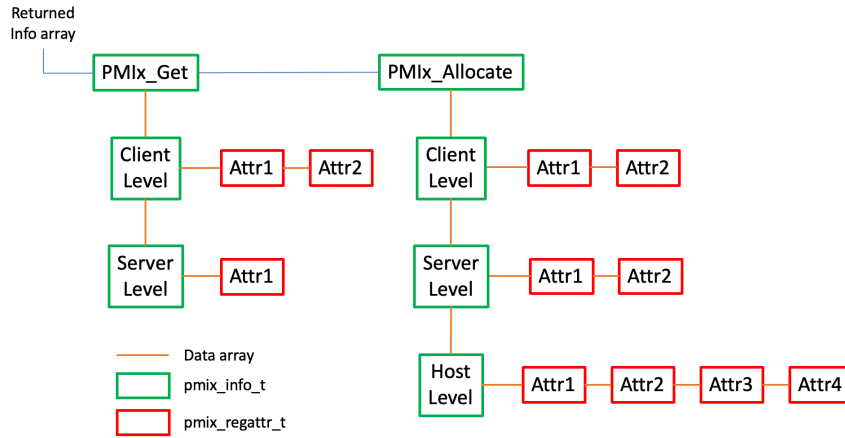


Figure 7.1.: Returned information hierarchy for attribute support request

3
4
5
6
7

The array of returned structures, and their child arrays, are subject to the return rules for the **PMIx_Query_info_nb** API. For example, a request for supported attributes of the **PMIx_Get** function that includes the *host* level will return values for the *client* and *server* levels, plus an array element with a *key* of **PMIX_HOST_ATTRIBUTES** and a value type of **PMIX_UNDEF** indicating that no attributes are supported at that level.

CHAPTER 8

Synchronization

1 Applications may need to synchronize their operations at various points in their execution.
2 Depending on a variety of factors (e.g., the programming model and where the synchronization
3 point lies), the application may choose to execute the operation using PMIx to access the
4 communication capabilities of the host environment's infrastructure. This is particularly useful in
5 situations where communication libraries are not yet initialized by the application. Synchronization
6 operations also offer an opportunity for processes to exchange data at a known point in their
7 execution. For example, communication libraries within different processes can synchronize to
8 exchange information on communication endpoints for subsequent wireup of messaging protocols.

9 PMIx clients can use the `PMIx_Fence` and `PMIx_Fence_nb` functions to synchronize a set of
10 processes. The fence operation can be useful after an application performs a number of `PMIx_Put`
11 operations to coordinate with other processes that the data is available for access. This avoids
12 unsuccessful `PMIx_Get` calls that might otherwise be invoked before the corresponding
13 `PMIx_Put` call is complete.

14 In its default form, the fence operation acts as a barrier between the processes and does not
15 exchange data. Clients can pass the `PMIX_COLLECT_DATA` attribute to request that the
16 `PMIx_Fence` and `PMIx_Fence_nb` functions exchange all committed data between all involved
17 servers during the synchronization operation. This will make local to each process the data put by
18 other processes resulting in faster resolution of `PMIx_Get` and `PMIx_Get_nb` function calls at
19 the cost of a synchronous data exchange and associated memory footprint expansion. In many
20 situations this attribute may have performance benefits as many systems are optimized for
21 transporting larger amounts of data. In such applications, a 'put/commit/fence/get' pattern is
22 common for efficiently exchanging key-value pairs. For applications where only a small subset of
23 clients access another small subset's key-value pairs this attribute may not be beneficial. As such,
24 applications are not required to use `PMIx_Fence` or `PMIx_Fence_nb` functions nor the
25 associated data collection attribute to ensure correctness of PMIx get/put functionality.

26 8.1 PMIx_Fence

27 Summary

28 Execute a blocking barrier across the processes identified in the specified array, collecting
29 information posted via `PMIx_Put` as directed.

Format

C

```
pmix_status_t
PMIx_Fence(const pmix_proc_t procs[], size_t nprocs,
           const pmix_info_t info[], size_t ninfo);
```

C

IN `procs`

Array of `pmix_proc_t` structures (array of handles)

IN `nprocs`

Number of elements in the `procs` array (integer)

IN `info`

Array of info structures (array of handles)

IN `ninfo`

Number of elements in the `info` array (integer)

Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_COLLECT_DATA "pmix.collect" (bool)

Collect all data posted by the participants using `PMIx_Put` that has been committed via `PMIx_Commit`, making the collection locally available to each participant at the end of the operation. By default, this will include all job-level information that was locally generated by PMIx servers unless excluded using the `PMIX_COLLECT_GENERATED_JOB_INFO` attribute.

PMIX_COLLECT_GENERATED_JOB_INFO "pmix.collect.gen" (bool)

Collect all job-level information (i.e., reserved keys) that was locally generated by PMIx servers. Some job-level information (e.g., distance between processes and fabric devices) is best determined on a distributed basis as it primarily pertains to local processes. Should remote processes need to access the information, it can either be obtained collectively using the `PMIx_Fence` operation with this directive, or can be retrieved one peer at a time using `PMIx_Get` without first having performed the job-wide collection.

Optional Attributes

The following attributes are optional for PMIx implementations:

PMIX_ALL_CLONES_PARTICIPATE "pmix.clone.part" (bool)

All *clones* of the calling process must participate in the collective operation.

The following attributes are optional for host environments:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Passing a **NULL** pointer as the *procs* parameter indicates that the fence is to span all processes in the client's namespace. Each provided `pmix_proc_t` struct can pass **PMIX_RANK_WILDCARD** to indicate that all processes in the given namespace are participating.

The *info* array is used to pass user directives regarding the behavior of the fence operation. Note that for scalability reasons, the default behavior for **PMIx_Fence** is to not collect data posted by the operation's participants.

Advice to PMIx library implementers

PMIx_Fence and its non-blocking form are both *collective* operations. Accordingly, the PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective.

8.2 PMIx_Fence_nb

Summary

Execute a nonblocking **PMIx_Fence** across the processes identified in the specified array of processes, collecting information posted via **PMIx_Put** as directed.

Format

C

```
pmix_status_t
PMIx_Fence_nb(const pmix_proc_t procs[], size_t nprocs,
              const pmix_info_t info[], size_t ninfo,
              pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

IN **procs**
Array of `pmix_proc_t` structures (array of handles)

IN **nprocs**
Number of elements in the *procs* array (integer)

IN **info**
Array of info structures (array of handles)

IN **ninfo**
Number of elements in the *info* array (integer)

IN **cbfunc**
Callback function (function reference)

IN **cbdata**
Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when `PMIX_SUCCESS` is returned.

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- `PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called. This can occur if the collective involved only processes on the local node.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

`PMIX_COLLECT_DATA` "pmix.collect" (bool)
Collect all data posted by the participants using `PMIx_Put` that has been committed via `PMIx_Commit`, making the collection locally available to each participant at the end of the operation. By default, this will include all job-level information that was locally generated by PMIx servers unless excluded using the `PMIX_COLLECT_GENERATED_JOB_INFO` attribute.

`PMIX_COLLECT_GENERATED_JOB_INFO` "pmix.collect.gen" (bool)

1 Collect all job-level information (i.e., reserved keys) that was locally generated by PMIx
2 servers. Some job-level information (e.g., distance between processes and fabric devices) is
3 best determined on a distributed basis as it primarily pertains to local processes. Should
4 remote processes need to access the information, it can either be obtained collectively using
5 the **PMIx_Fence** operation with this directive, or can be retrieved one peer at a time using
6 **PMIx_Get** without first having performed the job-wide collection.

Optional Attributes

7 The following attributes are optional for PMIx implementations:

8 **PMIX_ALL_CLONES_PARTICIPATE** "pmix.clone.part" (bool)

9 All *clones* of the calling process must participate in the collective operation.

10 The following attributes are optional for host environments that support this operation:

11 **PMIX_TIMEOUT** "pmix.timeout" (int)

12 Time in seconds before the specified operation should time out (zero indicating infinite) and
13 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
14 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

15 Nonblocking version of the **PMIx_Fence** routine. See the **PMIx_Fence** description for further
16 details.
17

18 8.2.1 Fence-related attributes

19 The following attributes are defined specifically to support the fence operation:

20 **PMIX_COLLECT_DATA** "pmix.collect" (bool)

21 Collect all data posted by the participants using **PMIx_Put** that has been committed via
22 **PMIx_Commit**, making the collection locally available to each participant at the end of the
23 operation. By default, this will include all job-level information that was locally generated
24 by PMIx servers unless excluded using the **PMIX_COLLECT_GENERATED_JOB_INFO**
25 attribute.

26 *Provisional*

PMIX_LOCAL_COLLECTIVE_STATUS "pmix.loc.col.st" (pmix_status_t)

27 Status code for local collective operation being reported to the host by the server library.
28 PMIx servers may aggregate the participation by local client processes in a collective
29 operation - e.g., instead of passing individual client calls to **PMIx_Fence** up to the host
30 environment, the server may pass only a single call to the host when all local participants
31 have executed their **PMIx_Fence** call, thereby reducing the burden placed on the host.
32 However, in cases where the operation locally fails (e.g., if a participating client abnormally
33 terminates prior to calling the operation), the server upcall functions to the host do not
34 include a **pmix_status_t** by which the PMIx server can alert the host to that failure.
35 This attribute resolves that problem by allowing the server to pass the status information
36 regarding the local collective operation.

Advice to PMIx server hosts

1 The PMIx server is allowed to pass **PMIX_SUCCESS** using this attribute, but is not required to do
2 so. PMIx implementations may choose to only report errors in this manner. The lack of an included
3 status shall therefore be taken to indicate that the collective operation locally succeeded.

4 **PMIX_COLLECT_GENERATED_JOB_INFO** "**pmix.collect.gen**" (bool)

5 Collect all job-level information (i.e., reserved keys) that was locally generated by PMIx
6 servers. Some job-level information (e.g., distance between processes and fabric devices) is
7 best determined on a distributed basis as it primarily pertains to local processes. Should
8 remote processes need to access the information, it can either be obtained collectively using
9 the **PMIx_Fence** operation with this directive, or can be retrieved one peer at a time using
10 **PMIx_Get** without first having performed the job-wide collection.

11 **PMIX_ALL_CLONES_PARTICIPATE** "**pmix.clone.part**" (bool)

12 All *clones* of the calling process must participate in the collective operation.

CHAPTER 9

Publish/Lookup Operations

Chapter 6 and Section 5.1 present how reserved and non-reserved keys deal with information that either is associated with a specific process (i.e., the retrieving process knows the identifier of the process that posted it) or requires a synchronization operation prior to retrieval (e.g., the case of globally unique non-reserved keys). However, another requirement exists for an asynchronous exchange of data where neither the posting nor the retrieving process is known in advance (e.g. two namespaces that do not share a child-parent relationship). The APIs defined in this section focus on resolving that specific situation by allowing processes to publish data that can subsequently be retrieved solely by referral to its key. Mechanisms for constraining the scope of availability of the information are also provided as a means for better targeting of the eventual recipient(s).

Note that no presumption is made regarding how the published information is to be stored, nor as to the entity (host environment or PMIx implementation) that shall act as the datastore. The descriptions in the remainder of this chapter shall simply refer to that entity as the *datastore*.

9.1 PMIx_Publish2

Summary

Publish data for later access via [PMIx_Lookup](#).

Format

Provisional
v5.0

```
pmix_status_t
```

```
PMIx_Publish2(const pmix_info_t pinfo[], size_t npinfo)
```

```
const pmix_info_t info[], size_t ninfo);
```

IN `pinfo`

Array of key value pairs to publish (array of [pmix_info_t](#))

IN `npinfo`

Number of elements in the *pinfo* array (integer)

IN `info`

Array of info structures (array of [pmix_info_t](#))

IN `ninfo`

Number of elements in the *info* array (integer)

Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Required Attributes

1 There are no required attributes for this API. PMIx implementations that do not directly support the
2 operation but are hosted by environments that do support it must pass any attributes that are
3 provided by the client to the host environment for processing. In addition, the PMIx library is
4 required to add the `PMIX_USERID` and the `PMIX_GRPID` attributes of the client process that
5 published the information to the *info* array passed to the host environment.

Optional Attributes

6 The following attributes are optional for host environments that support this operation:

7 **`PMIX_TIMEOUT`** "`pmix.timeout`" (`int`)

8 Time in seconds before the specified operation should time out (zero indicating infinite) and
9 return the `PMIX_ERR_TIMEOUT` error. Care should be taken to avoid race conditions
10 caused by multiple layers (client, server, and host) simultaneously timing the operation.

11 **`PMIX_RANGE`** "`pmix.range`" (`pmix_data_range_t`)

12 Define constraints on the processes that can access published data or generated events or
13 define constraints on the provider of data when looking up published data.

14 **`PMIX_PERSISTENCE`** "`pmix.persist`" (`pmix_persistence_t`)

15 Declare how long the datastore shall retain the provided data. The datastore is to delete the
16 data upon reaching the persistence criterion.

17 **`PMIX_ACCESS_USERIDS`** "`pmix.auids`" (`pmix_data_array_t`)

18 Array of effective User IDs (UIDs) that are allowed to access the published data.

19 **`PMIX_ACCESS_GRPIDS`** "`pmix.agids`" (`pmix_data_array_t`)

20 Array of effective Group IDs (GIDs) that are allowed to access the published data.

21 **`PMIX_ACCESS_PERMISSIONS`** "`pmix.aperms`" (`pmix_data_array_t`)

22 Define access permissions for the published data. The value shall contain an array of
23 `pmix_info_t` structs containing the specified permissions.

Description

Publish the data in the *pinfo* array for subsequent lookup. By default, the data will be published into the `PMIX_RANGE_SESSION` range and with `PMIX_PERSIST_APP` persistence. Changes to those values, and any additional directives, can be included in the *info* array. Attempts to access the data by processes outside of the provided data range shall be rejected. The `PMIX_PERSISTENCE` attribute instructs the datastore holding the published information as to how long that information is to be retained.

The blocking form of this call will block until it has obtained confirmation from the datastore that the data is available for lookup. The *info* array can be released upon return from the blocking function call.

Publishing duplicate keys is permitted provided they are published to different ranges. Custom ranges are considered different if they have different members. Duplicate keys being published on the same data range shall return the `PMIX_ERR_DUPLICATE_KEY` error.

9.2 PMIx_Publish2_nb

Summary

Nonblocking `PMIx_Publish2` routine.

Format

C

Provisional
v5.0

`pmix_status_t`

```
PMIx_Publish2_nb(const pmix_info_t pinfo[], size_t npinfo,  
                 const pmix_info_t info[], size_t ninfo,  
                 pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

- IN** `pinfo`
Array of key value pairs to publish (array of `pmix_info_t`)
- IN** `npinfo`
Number of elements in the *pinfo* array (integer)
- IN** `info`
Array of info structures (array of `pmix_info_t`)
- IN** `ninfo`
Number of elements in the *info* array (integer)
- IN** `cbfunc`
Callback function `pmix_op_cbfunc_t` (function reference)
- IN** `cbdata`
Data to be passed to the callback function (memory reference)

1 A successful return indicates that the request is being processed and the result will be returned in
2 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
3 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

4 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 5 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
6 returned *success* - the *cbfunc* will *not* be called.

7 If none of the above return codes are appropriate, then an implementation must return either a
8 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

9 There are no required attributes for this API. PMIx implementations that do not directly support the
10 operation but are hosted by environments that do support it must pass any attributes that are
11 provided by the client to the host environment for processing. In addition, the PMIx library is
12 required to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that
13 published the information to the *info* array passed to the host environment.

Optional Attributes

14 The following attributes are optional for host environments that support this operation:

15 **PMIX_TIMEOUT** "**pmix.timeout**" (**int**)

16 Time in seconds before the specified operation should time out (zero indicating infinite) and
17 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
18 caused by multiple layers (client, server, and host) simultaneously timing the operation.

19 **PMIX_RANGE** "**pmix.range**" (**pmix_data_range_t**)

20 Define constraints on the processes that can access published data or generated events or
21 define constraints on the provider of data when looking up published data.

22 **PMIX_PERSISTENCE** "**pmix.persist**" (**pmix_persistence_t**)

23 Declare how long the datastore shall retain the provided data. The datastore is to delete the
24 data upon reaching the persistence criterion.

25 **PMIX_ACCESS_USERIDS** "**pmix.auids**" (**pmix_data_array_t**)

26 Array of effective UIDs that are allowed to access the published data.

27 **PMIX_ACCESS_GRPIDS** "**pmix.agids**" (**pmix_data_array_t**)

28 Array of effective GIDs that are allowed to access the published data.

29 **PMIX_ACCESS_PERMISSIONS** "**pmix.aperms**" (**pmix_data_array_t**)

30 Define access permissions for the published data. The value shall contain an array of
31 **pmix_info_t** structs containing the specified permissions.

Description

Nonblocking `PMIx_Publish2` routine.

9.3 PMIx_Publish

Summary

Publish data for later access via `PMIx_Lookup`.

Format

PMIx v1.0

```
pmix_status_t  
PMIx_Publish(const pmix_info_t info[], size_t ninfo);
```

IN info

Array of info structures containing both data to be published and directives (array of handles)

IN ninfo

Number of elements in the *info* array (integer)

Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

There are no required attributes for this API. PMIx implementations that do not directly support the operation but are hosted by environments that do support it must pass any attributes that are provided by the client to the host environment for processing. In addition, the PMIx library is required to add the `PMIX_USERID` and the `PMIX_GRPID` attributes of the client process that published the information to the *info* array passed to the host environment.

Optional Attributes

The following attributes are optional for host environments that support this operation:

`PMIX_TIMEOUT` "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the `PMIX_ERR_TIMEOUT` error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

`PMIX_RANGE` "pmix.range" (pmix_data_range_t)

Define constraints on the processes that can access published data or generated events or define constraints on the provider of data when looking up published data.

`PMIX_PERSISTENCE` "pmix.persist" (pmix_persistence_t)

Declare how long the datastore shall retain the provided data. The datastore is to delete the data upon reaching the persistence criterion.

`PMIX_ACCESS_USERIDS` "pmix.auids" (pmix_data_array_t)

1 Array of effective UIDs that are allowed to access the published data.
 2 **PMIX_ACCESS_GRPIDS** "pmix.agids" (**pmix_data_array_t**)
 3 Array of effective GIDs that are allowed to access the published data.
 4 **PMIX_ACCESS_PERMISSIONS** "pmix.aperms" (**pmix_data_array_t**)
 5 Define access permissions for the published data. The value shall contain an array of
 6 **pmix_info_t** structs containing the specified permissions.



7 **Description**

8 Publish the data in the *info* array for subsequent lookup. By default, the data will be published into
 9 the **PMIX_RANGE_SESSION** range and with **PMIX_PERSIST_APP** persistence. Changes to
 10 those values, and any additional directives, can be included in the **pmix_info_t** array. Attempts
 11 to access the data by processes outside of the provided data range shall be rejected. The
 12 **PMIX_PERSISTENCE** attribute instructs the datastore holding the published information as to
 13 how long that information is to be retained.

14 The blocking form of this call will block until it has obtained confirmation from the datastore that
 15 the data is available for lookup. The *info* array can be released upon return from the blocking
 16 function call.

17 Publishing duplicate keys is permitted provided they are published to different ranges. Custom
 18 ranges are considered different if they have different members. Duplicate keys being published on
 19 the same data range shall return the **PMIX_ERR_DUPLICATE_KEY** error.

20 **9.4 PMIx_Publish_nb**

21 **Summary**

22 Nonblocking **PMIx_Publish** routine.

23 **Format**

PMIx v1.0

24 **pmix_status_t**
 25 **PMIx_Publish_nb**(const **pmix_info_t** info[], **size_t** ninfo,
 26 **pmix_op_cbfunc_t** cbfunc, void *cbdata);

27 **IN** **info**

28 Array of info structures containing both data to be published and directives (array of handles)

29 **IN** **ninfo**

30 Number of elements in the *info* array (integer)

31 **IN** **cbfunc**

32 Callback function **pmix_op_cbfunc_t** (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning from the API. The callback function, `cbfunc`, is only called when `PMIX_SUCCESS` is returned.

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- `PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed and returned *success* - the `cbfunc` will *not* be called.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

There are no required attributes for this API. PMIx implementations that do not directly support the operation but are hosted by environments that do support it must pass any attributes that are provided by the client to the host environment for processing. In addition, the PMIx library is required to add the `PMIX_USERID` and the `PMIX_GRPID` attributes of the client process that published the information to the `info` array passed to the host environment.

Optional Attributes

The following attributes are optional for host environments that support this operation:

`PMIX_TIMEOUT` "`pmix.timeout`" (`int`)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the `PMIX_ERR_TIMEOUT` error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

`PMIX_RANGE` "`pmix.range`" (`pmix_data_range_t`)

Define constraints on the processes that can access published data or generated events or define constraints on the provider of data when looking up published data.

`PMIX_PERSISTENCE` "`pmix.persist`" (`pmix_persistence_t`)

Declare how long the datastore shall retain the provided data. The datastore is to delete the data upon reaching the persistence criterion.

`PMIX_ACCESS_USERIDS` "`pmix.auids`" (`pmix_data_array_t`)

Array of effective UIDs that are allowed to access the published data.

`PMIX_ACCESS_GRPIDS` "`pmix.agids`" (`pmix_data_array_t`)

Array of effective GIDs that are allowed to access the published data.

`PMIX_ACCESS_PERMISSIONS` "`pmix.aperms`" (`pmix_data_array_t`)

Define access permissions for the published data. The value shall contain an array of `pmix_info_t` structs containing the specified permissions.

1 **Description**

2 Nonblocking `PMIx_Publish` routine.

3 **9.5 Publish-specific constants**

4 The following constants are defined for use with the `PMIx_Publish` APIs:

5 `PMIX_ERR_DUPLICATE_KEY` -53 The provided key has already been published on the
6 same data range.

7 **9.6 Publish-specific attributes**

8 The following attributes are defined for use with the `PMIx_Publish` APIs:

9 `PMIX_RANGE` "pmix.range" (`pmix_data_range_t`)

10 Define constraints on the processes that can access published data or generated events or
11 define constraints on the provider of data when looking up published data.

12 `PMIX_PERSISTENCE` "pmix.persist" (`pmix_persistence_t`)

13 Declare how long the datastore shall retain the provided data. The datastore is to delete the
14 data upon reaching the persistence criterion.

15 `PMIX_ACCESS_PERMISSIONS` "pmix.aperms" (`pmix_data_array_t`)

16 Define access permissions for the published data. The value shall contain an array of
17 `pmix_info_t` structs containing the specified permissions.

18 `PMIX_ACCESS_USERIDS` "pmix.auids" (`pmix_data_array_t`)

19 Array of effective UIDs that are allowed to access the published data.

20 `PMIX_ACCESS_GRPIDS` "pmix.agids" (`pmix_data_array_t`)

21 Array of effective GIDs that are allowed to access the published data.

22 **9.7 Publish-Lookup Datatypes**

23 The following data types are defined for use with the `PMIx_Publish` APIs.

1 9.7.1 Range of Published Data

2 The `pmix_data_range_t` structure is a `uint8_t` type that defines a range for data *published*
3 via the `PMIx_Publish` API and events generated via the `PMIx_Notify_event`. The
4 following constants can be used to set a variable of the type `pmix_data_range_t`.

5 **PMIX_RANGE_UNDEF 0** Undefined range.

6 **PMIX_RANGE_RM 1** Data is intended for the host environment, or lookup is restricted to data
7 published by the host environment.

8 **PMIX_RANGE_LOCAL 2** Published data and generated events are restricted to processes on
9 the same node as the publisher or event creator. Lookup of data is restricted to data published
10 by processes on the same node as the requester.

11 **PMIX_RANGE_NAMESPACE 3** Published data and generated events are restricted to
12 processes in the same namespace as the publisher or event creator. Lookup of data is
13 restricted to data published by processes in the same namespace as the requester.

14 **PMIX_RANGE_SESSION 4** Published data and generated events are restricted to processes in
15 the same session as the publisher or event creator. Lookup of data is restricted to data
16 published by processes in the same session as the requester.

17 **PMIX_RANGE_GLOBAL 5** Published data and generated events are available to all processes
18 within the domain of the host environment. Lookup of data is unrestricted and open to data
19 published by any processes within the domain of the host environment as the requester. This
20 range differs from `PMIX_RANGE_RM` only on systems which have mechanisms to share
21 events and publish/lookup data across multiple instances of a host environment.

22 **PMIX_RANGE_PROC_LOCAL 7** Published data and generated events are available only to
23 calling process. Lookup of data is restricted to data published by the calling process.

24 **PMIX_RANGE_CUSTOM 6** Published data and generated events are restricted to processes
25 described in the `pmix_info_t` associated with this call. Lookup of data is restricted to data
26 published by the processes described in in the `pmix_info_t`.

27 **PMIX_RANGE_INVALID UINT8_MAX** Invalid value - typically used to indicate that a range
28 has not yet been set.

29 9.7.2 Data Persistence Structure

30 *PMIx v1.0*

31 The `pmix_persistence_t` structure is a `uint8_t` type that defines the policy for data
32 published by clients via the `PMIx_Publish` API. The following constants can be used to set a
variable of the type `pmix_persistence_t`.

33 **PMIX_PERSIST_INDEF 0** Retain data until unpublished.

34 **PMIX_PERSIST_FIRST_READ 1** Retain data until the first access, then the data is deleted.

35 **PMIX_PERSIST_PROC 2** Retain data until the publishing process terminates.

36 **PMIX_PERSIST_APP 3** Retain data until the application terminates.

37 **PMIX_PERSIST_SESSION 4** Retain data until the session/allocation terminates.

38 **PMIX_PERSIST_INVALID UINT8_MAX** Invalid value - typically used to indicate that a
39 persistence has not yet been set.

1 9.7.3 Lookup Related Data Structures

2 The `pmix_pdata_t` structure is used both to request the lookup of keys and to describe the value
3 and publishing process of any keys that were successfully retrieved. A request to lookup published
4 values is described by an array of `pmix_pdata_t` structures. Only the key field is used in the
5 lookup request. The results of the lookup operation are returned in the same array with the proc and
6 value fields set when the key is successfully found. The value field's data type is set to
7 `PMIX_UNDEF` in the associated *value* struct of any key which was not retrieved.

```
8 typedef struct pmix_pdata {  
9     pmix_proc_t proc;  
10    pmix_key_t key;  
11    pmix_value_t value;  
12 } pmix_pdata_t;
```

13 where:

- 14 • *proc* is the process identifier of the data publisher.
- 15 • *key* is the string key of the published data.
- 16 • *value* is the value associated with the *key*.

17 9.8 PMIx_Lookup

18 Summary

19 Lookup information published by a process or host environment using `PMIx_Publish` or
20 `PMIx_Publish_nb`.

21 Format

PMIx v1.0

```
22 pmix_status_t  
23 PMIx_Lookup(pmix_pdata_t data[], size_t ndata,  
24             const pmix_info_t info[], size_t ninfo);
```

25 INOUT data

26 Array of publishable data structures (array of `pmix_pdata_t`)

27 IN ndata

28 Number of elements in the *data* array (integer)

29 IN info

30 Array of info structures (array of `pmix_info_t`)

1 **IN ninfo**

2 Number of elements in the *info* array (integer)

3 Returns PMIX_SUCCESS or one of the following error codes when the condition described occurs:

- 4 ● **PMIX_ERR_NOT_FOUND** None of the requested data could be found within the requester's
5 range.
- 6 ● **PMIX_ERR_PARTIAL_SUCCESS** Some of the requested data was found. Any key that cannot
7 be found will return with a data type of **PMIX_UNDEF** in the associated *value* struct. Note that
8 the specific reason for a particular piece of missing information (e.g., lack of permissions) cannot
9 be communicated back to the requester in this situation.
- 10 ● **PMIX_ERR_NO_PERMISSIONS** All requested data was found and range restrictions were met
11 for each specified key, but none of the matching data could be returned due to lack of access
12 permissions.

13 If none of the above return codes are appropriate, then an implementation must return either a
14 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

▼----- Required Attributes -----▼

15 PMIx libraries are not required to directly support any attributes for this function. However, any
16 provided attributes must be passed to the host environment for processing, and the PMIx library is
17 required to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that is
18 requesting the info.

▲----- Optional Attributes -----▲

19 The following attributes are optional for host environments that support this operation:

20 **PMIX_TIMEOUT** "**pmix.timeout**" (**int**)

21 Time in seconds before the specified operation should time out (zero indicating infinite) and
22 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
23 caused by multiple layers (client, server, and host) simultaneously timing the operation.

24 **PMIX_RANGE** "**pmix.range**" (**pmix_data_range_t**)

25 Define constraints on the processes that can access published data or generated events or
26 define constraints on the provider of data when looking up published data.

27 **PMIX_WAIT** "**pmix.wait**" (**int**)

28 Caller requests that the PMIx server wait until at least the specified number of values are
29 found (a value of zero indicates *all* and is the default).

Description

Lookup information published by a process or host environment using `PMIx_Publish` or `PMIx_Publish_nb`. A lookup operation is always performed on a range which can be specified using the directive `PMIX_RANGE` or otherwise defaults to `PMIX_RANGE_SESSION`.

The lookup operation will be constrained to data published to the specified range. Data is returned per the retrieval rules of Section 9.10.

The *data* parameter consists of an array of `pmix_pdata_t` structures with the keys specifying the requested information. Data will be returned for each **key** field in the associated **value** field of this structure as per the above description of return values. The **proc** field in each `pmix_pdata_t` structure will contain the namespace/rank of the process that published the data.

Advice to users

Although this is a blocking function, it will not wait by default for the requested data to be published. Instead, it will block for the time required by the datastore to lookup its current data and return any found items. Thus, the caller is responsible for either ensuring that data is published prior to executing a lookup, using `PMIX_WAIT` to instruct the datastore to wait for the data to be published, or retrying until the requested data is found.

9.9 PMIx_Lookup_nb

Summary

Nonblocking version of `PMIx_Lookup`.

Format

PMIx v1.0

```
pmix_status_t
PMIx_Lookup_nb(char **keys,
               const pmix_info_t info[], size_t ninfo,
               pmix_lookup_cbfunc_t cbfunc, void *cbdata);
```

IN	keys	NULL-terminated array of keys (array of strings)
IN	info	Array of info structures (array of handles)
IN	ninfo	Number of elements in the <i>info</i> array (integer)
IN	cbfunc	Callback function (handle)

IN `cbdata`

Callback data to be provided to the callback function (pointer)

A successful return indicates that the request is being processed and the result will be returned in the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning from the API. The callback function, `cbfunc`, is only called when `PMIX_SUCCESS` is returned.

If executed, the status returned in the provided callback function will be one of the following constants:

- `PMIX_SUCCESS` All data was found and has been returned.
- `PMIX_ERR_NOT_FOUND` None of the requested data was available within the requester's range. The `pdata` array in the callback function shall be `NULL` and the `npdata` parameter set to zero.
- `PMIX_ERR_PARTIAL_SUCCESS` Some of the requested data was found. Only found data will be included in the returned `pdata` array. Note that the specific reason for a particular piece of missing information (e.g., lack of permissions or the data has not been published) cannot be communicated back to the requester in this situation.
- `PMIX_ERR_NOT_SUPPORTED` There is no available datastore (either at the host environment or PMIx implementation level) on this system that supports this function.
- `PMIX_ERR_NO_PERMISSIONS` All of the requested data was found and range restrictions were met for each specified key, but none of the matching data could be returned due to lack of access permissions.
- a non-zero PMIx error constant indicating a reason for the request's failure.

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host environment for processing, and the PMIx library is required to add the `PMIX_USERID` and the `PMIX_GRPID` attributes of the client process that is requesting the info.

Optional Attributes

The following attributes are optional for host environments that support this operation:

`PMIX_TIMEOUT` "`pmix.timeout`" (`int`)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the `PMIX_ERR_TIMEOUT` error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

`PMIX_RANGE` "`pmix.range`" (`pmix_data_range_t`)

Define constraints on the processes that can access published data or generated events or define constraints on the provider of data when looking up published data.

`PMIX_WAIT` "`pmix.wait`" (`int`)

1 Caller requests that the PMIx server wait until at least the specified number of values are
2 found (a value of zero indicates *all* and is the default).



3 Description

4 Non-blocking form of the [PMIx_Lookup](#) function.

5 9.9.0.1 Lookup data structure support macros

6 The following macros are provided to support the [pmix_pdata_t](#) structure.

7 **Static initializer for the pdata structure**

8 **Provisional** Provide a static initializer for the [pmix_pdata_t](#) fields.

PMIx v4.2

▼ C _____ ▼

9 **PMIX_LOOKUP_STATIC_INIT**

▲ _____ C _____ ▲

10 **Initialize the pdata structure**

11 Initialize the [pmix_pdata_t](#) fields

PMIx v1.0

▼ C _____ ▼

12 **PMIX_PDATA_CONSTRUCT (m)**

▲ _____ C _____ ▲

13 **IN m**

14 Pointer to the structure to be initialized (pointer to [pmix_pdata_t](#))

15 **Destruct the pdata structure**

16 Destruct the [pmix_pdata_t](#) fields

PMIx v1.0

▼ C _____ ▼

17 **PMIX_PDATA_DESTRUCT (m)**

▲ _____ C _____ ▲

18 **IN m**

19 Pointer to the structure to be destructed (pointer to [pmix_pdata_t](#))

20 **Create a pdata array**

21 Allocate and initialize an array of [pmix_pdata_t](#) structures

PMIx v1.0

▼ C _____ ▼

22 **PMIX_PDATA_CREATE (m, n)**

▲ _____ C _____ ▲

23 **INOUT m**

24 Address where the pointer to the array of [pmix_pdata_t](#) structures shall be stored (handle)

25 **IN n**

26 Number of structures to be allocated ([size_t](#))

1 **Free a pdata structure**
2 Release a `pmix_pdata_t` structure

PMIx v4.0

▼ `PMIX_PDATALOAD` C

3 **PMIX_PDATALOAD (m)**

▲ `PMIX_PDATALOAD` C

4 **IN** `m`
5 Pointer to a `pmix_pdata_t` structure (handle)

6 **Free a pdata array**
7 Release an array of `pmix_pdata_t` structures

PMIx v1.0

▼ `PMIX_PDATALOAD` C

8 **PMIX_PDATALOAD (m, n)**

▲ `PMIX_PDATALOAD` C

9 **IN** `m`
10 Pointer to the array of `pmix_pdata_t` structures (handle)

11 **IN** `n`
12 Number of structures in the array (`size_t`)

13 **Load a lookup data structure**
14 This macro simplifies the loading of key, process identifier, and data into a `pmix_pdata_t` by
15 correctly assigning values to the structure's fields.

PMIx v1.0

▼ `PMIX_PDATALOAD` C

16 **PMIX_PDATALOAD (m, p, k, d, t);**

▲ `PMIX_PDATALOAD` C

17 **IN** `m`
18 Pointer to the `pmix_pdata_t` structure into which the key and data are to be loaded
19 (pointer to `pmix_pdata_t`)

20 **IN** `p`
21 Pointer to the `pmix_proc_t` structure containing the identifier of the process being
22 referenced (pointer to `pmix_proc_t`)

23 **IN** `k`
24 String key to be loaded - must be less than or equal to `PMIX_MAX_KEYLEN` in length
25 (handle)

26 **IN** `d`
27 Pointer to the data value to be loaded (handle)

28 **IN** `t`
29 Type of the provided data value (`pmix_data_type_t`)

Advice to users

1 Key, process identifier, and data will all be copied into the `pmix_pdata_t` - thus, the source
2 information can be modified or free'd without affecting the copied data once the macro has
3 completed.

Transfer a lookup data structure

4 This macro simplifies the transfer of key, process identifier, and data value between
5 two `pmix_pdata_t` structures.
6

PMIx v2.0

```
7 PMIX_PDATA_XFER(d, s);
```

8 **IN** d

9 Pointer to the destination `pmix_pdata_t` (pointer to `pmix_pdata_t`)

10 **IN** s

11 Pointer to the source `pmix_pdata_t` (pointer to `pmix_pdata_t`)

Advice to users

12 Key, process identifier, and data will all be copied into the destination `pmix_pdata_t` - thus, the
13 source `pmix_pdata_t` may free'd without affecting the copied data once the macro has
14 completed.

1 9.10 Retrieval rules for published data

2 The retrieval rules for published data primarily revolve around enforcing data access permissions
3 and range constraints. All publish and lookup operations operate on a range. If not specified, the
4 range defaults to **PMIX_RANGE_SESSION**. The key being looked up will match with a published
5 key only if all of the following conditions are met:

- 6 1. The lookup key matches the published key.
- 7 2. The type of range specified by the publisher is the same as the type of range specified by the
8 requester.
- 9 3. The requestor must be a member of the range specified by the publisher.
- 10 4. The publisher must be a member of the range specified by the requestor.
- 11 5. If the publisher specified access permissions, the effective UID and GID of the requester must
12 meet those requirements.

13 The status returned by the datastore shall be set to:

- 14 • **PMIX_SUCCESS** All data was found and is included in the returned information.
- 15 • **PMIX_ERR_NOT_FOUND** None of the requested data could be found within a requester's range.
- 16 • **PMIX_ERR_PARTIAL_SUCCESS** Some of the requested data was found. Only found data will
17 be included in the returned information. Note that the specific reason for a particular piece of
18 missing information (e.g., lack of permissions) cannot be communicated back to the requester in
19 this situation.
- 20 • **PMIX_ERR_NO_PERMISSIONS** All requested data was found and range restrictions were met
21 for each specified key, but none of the matching data could be returned due to lack of access
22 permissions.
- 23 • a non-zero PMIx error constant indicating a reason for the request's failure.

24 9.11 PMIx_Unpublish

25 Summary

26 Unpublish a list of keys published by the calling process.

27 Format

PMIx v1.0

C

```
28 pmix_status_t  
29 PMIx_Unpublish(char **keys,  
30                const pmix_info_t info[], size_t ninfo);
```

1 **IN** **keys**
 2 **NULL**-terminated array of keys (array of strings)
 3 **IN** **info**
 4 Array of info structures (array of handles)
 5 **IN** **ninfo**
 6 Number of elements in the *info* array (integer)

7 Returns **PMIX_SUCCESS** or a negative value indicating the error.

Required Attributes

8 PMIx libraries are not required to directly support any attributes for this function. However, any
 9 provided attributes must be passed to the host environment for processing, and the PMIx library is
 10 required to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process that is
 11 requesting the operation.

Optional Attributes

12 The following attributes are optional for host environments that support this operation:

13 **PMIX_TIMEOUT** "**pmix.timeout**" (**int**)
 14 Time in seconds before the specified operation should time out (zero indicating infinite) and
 15 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
 16 caused by multiple layers (client, server, and host) simultaneously timing the operation.
 17 **PMIX_RANGE** "**pmix.range**" (**pmix_data_range_t**)
 18 Define constraints on the processes that can access published data or generated events or
 19 define constraints on the provider of data when looking up published data.

Description

21 Unpublish a list of keys published by the calling process. The function will block until the data has
 22 been removed by the server (i.e., it is safe to publish that key again within the specified range). A
 23 value of **NULL** for the *keys* parameter instructs the server to remove all data published by this
 24 process.

25 By default, the range is assumed to be **PMIX_RANGE_SESSION**. Changes to the range, and any
 26 additional directives, can be provided in the *info* array.

27 9.12 PMIx_Unpublish_nb

28 Summary

29 Nonblocking version of **PMIx_Unpublish**.

Format

C

```
pmix_status_t
PMIx_Unpublish_nb(char **keys,
                  const pmix_info_t info[], size_t ninfo,
                  pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

IN keys

NULL-terminated array of keys (array of strings)

IN info

Array of info structures (array of handles)

IN ninfo

Number of elements in the *info* array (integer)

IN cbfunc

Callback function [pmix_op_cbfunc_t](#) (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when [PMIX_SUCCESS](#) is returned.

Returns [PMIX_SUCCESS](#) or one of the following error codes when the condition described occurs:

- [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host environment for processing, and the PMIx library is required to add the [PMIX_USERID](#) and the [PMIX_GRPID](#) attributes of the client process that is requesting the operation.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Define constraints on the processes that can access published data or generated events or define constraints on the provider of data when looking up published data.

Description

Non-blocking form of the **PMIx_Unpublish** function. The callback function will be executed once the server confirms removal of the specified data. The *info* array must be maintained until the callback is provided.

CHAPTER 10

Event Notification

1 This chapter defines the PMIx event notification system. These interfaces are designed to support
2 the reporting of events to/from clients and servers, and between library layers within a single
3 process.

4 10.1 Notification and Management

5 PMIx event notification provides an asynchronous out-of-band mechanism for communicating
6 events between application processes and/or elements of the SMS. Its uses span a wide range
7 including fault notification, coordination between multiple programming libraries within a single
8 process, and workflow orchestration for non-synchronous programming models. Events can be
9 divided into two distinct classes:

- 10 • *Job-specific events* directly relate to a job executing within the session, such as a debugger
11 attachment, process failure within a related job, or events generated by an application process.
12 Events in this category are to be immediately delivered to the PMIx server library for relay to the
13 related local processes.
- 14 • *Environment events* indirectly relate to a job but do not specifically target the job itself. This
15 category includes SMS-generated events such as Error Check and Correction (ECC) errors,
16 temperature excursions, and other non-job conditions that might directly affect a session's
17 resources, but would never include an event generated by an application process. Note that
18 although these do potentially impact the session's jobs, they are not directly tied to those jobs.
19 Thus, events in this category are to be delivered to the PMIx server library only upon request.

20 Both SMS elements and applications can register for events of either type.

▼ Advice to PMIx library implementers ▼

21 Race conditions can cause the registration to come after events of possible interest (e.g., a memory
22 ECC event that occurs after start of execution but prior to registration, or an application process
23 generating an event prior to another process registering to receive it). SMS vendors are *requested* to
24 cache environment events for some time to mitigate this situation, but are not *required* to do so.
25 However, PMIx implementers are *required* to cache all events received by the PMIx server library
26 and to deliver them to registering clients in the same order in which they were received

Advice to users

1 Applications must be aware that they may not receive environment events that occur prior to
2 registration, depending upon the capabilities of the host SMS.

3 The generator of an event can specify the *target range* for delivery of that event. Thus, the generator
4 can choose to limit notification to processes on the local node, processes within the same job as the
5 generator, processes within the same allocation, other threads within the same process, only the
6 SMS (i.e., not to any application processes), all application processes, or to a custom range based
7 on specific process identifiers. Only processes within the given range that register for the provided
8 event code will be notified. In addition, the generator can use attributes to direct that the event not
9 be delivered to any default event handlers, or to any multi-code handler (as defined below).

10 Event notifications provide the process identifier of the source of the event plus the event code and
11 any additional information provided by the generator. When an event notification is received by a
12 process, the registered handlers are scanned for their event code(s), with matching handlers
13 assembled into an *event chain* for servicing. Note that users can also specify a *source range* when
14 registering an event (using the same range designators described above) to further limit when they
15 are to be invoked. When assembled, PMIx event chains are ordered based on both the specificity of
16 the event handler and user directives at time of handler registration. By default, handlers are
17 grouped into three categories based on the number of event codes that can trigger the callback:

- 18 ● *single-code* handlers are serviced first as they are the most specific. These are handlers that are
19 registered against one specific event code.
- 20 ● *multi-code* handlers are serviced once all single-code handlers have completed. The handler will
21 be included in the chain upon receipt of an event matching any of the provided codes.
- 22 ● *default* handlers are serviced once all multi-code handlers have completed. These handlers are
23 always included in the chain unless the generator specifically excludes them.

24 Users can specify the callback order of a handler within its category at the time of registration.
25 Ordering can be specified by providing the relevant event handler names, if the user specified an
26 event handler name when registering the corresponding event. Thus, users can specify that a given
27 handler be executed before or after another handler should both handlers appear in an event chain
28 (the ordering is ignored if the other handler isn't included). Note that ordering does not imply
29 immediate relationships. For example, multiple handlers registered to be serviced after event
30 handler *A* will all be executed after *A*, but are not guaranteed to be executed in any particular order
31 amongst themselves.

32 In addition, one event handler can be declared as the *first* handler to be executed in the chain. This
33 handler will *always* be called prior to any other handler, regardless of category, provided the
34 incoming event matches both the specified range and event code. Only one handler can be so
35 designated — attempts to designate additional handlers as *first* will return an error. Deregistration
36 of the declared *first* handler will re-open the position for subsequent assignment.

1 Similarly, one event handler can be declared as the *last* handler to be executed in the chain. This
2 handler will *always* be called after all other handlers have executed, regardless of category,
3 provided the incoming event matches both the specified range and event code. Note that this
4 handler will not be called if the chain is terminated by an earlier handler. Only one handler can be
5 designated as *last* — attempts to designate additional handlers as *last* will return an error.
6 Deregistration of the declared *last* handler will re-open the position for subsequent assignment.

Advice to users

7 Note that the *last* handler is called *after* all registered default handlers that match the specified
8 range of the incoming event unless a handler prior to it terminates the chain. Thus, if the application
9 intends to define a *last* handler, it should ensure that no default handler aborts the process before it.

10 Upon completing its work and prior to returning, each handler *must* call the event handler
11 completion function provided when it was invoked (including a status code plus any information to
12 be passed to later handlers) so that the chain can continue being progressed. PMIx automatically
13 aggregates the status and any results of each handler (as provided in the completion callback) with
14 status from all prior handlers so that each step in the chain has full knowledge of what preceded it.
15 An event handler can terminate all further progress along the chain by passing the
16 [PMIX_EVENT_ACTION_COMPLETE](#) status to the completion callback function.

10.1.1 Events versus status constants

18 Return status constants (see Section [3.1.1](#)) represent values that can be returned from or passed into
19 PMIx APIs. These are distinct from PMIx *events* in that they are not values that can be registered
20 against event handlers. In general, the two types of constants are distinguished by inclusion of an
21 "ERR" in the name of error constants versus an "EVENT" in events, though there are exceptions
22 (e.g, the [PMIX_SUCCESS](#) constant).

10.1.2 PMIx_Register_event_handler

Summary

Register an event handler.

Format

C

```
pmix_status_t
PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,
                           pmix_info_t info[], size_t ninfo,
                           pmix_notification_fn_t evhdlr,
                           pmix_hdlr_reg_cbfunc_t cbfunc,
                           void *cbdata);
```

C

IN codes

Array of status codes (array of `pmix_status_t`)

IN ncodes

Number of elements in the *codes* array (`size_t`)

IN info

Array of info structures (array of handles)

IN ninfo

Number of elements in the *info* array (`size_t`)

IN evhdlr

Event handler to be called `pmix_notification_fn_t` (function reference)

IN cbfunc

Callback function `pmix_hdlr_reg_cbfunc_t` (function reference)

IN cbdata

Data to be passed to the *cbfunc* callback function (memory reference)

If *cbfunc* is `NULL`, the function call will be treated as a *blocking* call. In this case, the returned status will be either (a) the event handler reference identifier if the value is greater than or equal to zero, or (b) a negative error code indicative of the reason for the failure.

If the *cbfunc* is non-`NULL`, the function call will be treated as a *non-blocking* call and will return the following:

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when `PMIX_SUCCESS` is returned. The result of the registration operation shall be returned in the provided callback function along with the assigned event handler identifier.

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- `PMIX_ERR_EVENT_REGISTRATION` indicating that the registration has failed for an undetermined reason.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

1 The callback function must not be executed prior to returning from the API, and no events
2 corresponding to this registration may be delivered prior to the completion of the registration
3 callback function (*cbfunc*).

Required Attributes

4 The following attributes are required to be supported by all PMIx libraries:

5 **PMIX_EVENT_HDLR_NAME** "pmix.evname" (char*)

6 String name identifying this handler.

7 **PMIX_EVENT_HDLR_FIRST** "pmix.evfirst" (bool)

8 Invoke this event handler before any other handlers.

9 **PMIX_EVENT_HDLR_LAST** "pmix.evlast" (bool)

10 Invoke this event handler after all other handlers have been called.

11 **PMIX_EVENT_HDLR_FIRST_IN_CATEGORY** "pmix.evfirstcat" (bool)

12 Invoke this event handler before any other handlers in this category.

13 **PMIX_EVENT_HDLR_LAST_IN_CATEGORY** "pmix.evlastcat" (bool)

14 Invoke this event handler after all other handlers in this category have been called.

15 **PMIX_EVENT_HDLR_BEFORE** "pmix.evbefore" (char*)

16 Put this event handler immediately before the one specified in the (char*) value.

17 **PMIX_EVENT_HDLR_AFTER** "pmix.evafter" (char*)

18 Put this event handler immediately after the one specified in the (char*) value.

19 **PMIX_EVENT_HDLR_PREPEND** "pmix.evprepend" (bool)

20 Prepend this handler to the precedence list within its category.

21 **PMIX_EVENT_HDLR_APPEND** "pmix.evappend" (bool)

22 Append this handler to the precedence list within its category.

23 **PMIX_EVENT_CUSTOM_RANGE** "pmix.evrangle" (pmix_data_array_t*)

24 Array of **pmix_proc_t** defining range of event notification.

25 **PMIX_RANGE** "pmix.range" (pmix_data_range_t)

26 Define constraints on the processes that can access published data or generated events or
27 define constraints on the provider of data when looking up published data.

28 **PMIX_EVENT_RETURN_OBJECT** "pmix.evobject" (void *)

29 Object to be returned whenever the registered callback function **cbfunc** is invoked. The
30 object will only be returned to the process that registered it.

1
2 Host environments that implement support for PMIx event notification are required to support the
3 following attributes when registering handlers - these attributes are used to direct that the handler
4 should be invoked only when the event affects the indicated process(es):

5 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (pmix_proc_t)

6 The single process that was affected.

7 **PMIX_EVENT_AFFECTED_PROCS** "pmix.evaffected" (pmix_data_array_t*)

8 Array of pmix_proc_t defining affected processes.

9 **Description**

10 Register an event handler to report events. Note that the codes being registered do *not* need to be
11 PMIx error constants — any integer value can be registered. This allows for registration of
12 non-PMIx events such as those defined by a particular SMS vendor or by an application itself.

13 **Advice to users**

14 In order to avoid potential conflicts, users are advised to only define codes that lie outside the range
15 of the PMIx standard's error codes. Thus, SMS vendors and application developers should
16 constrain their definitions to positive values or negative values beyond the
PMIX_EXTERNAL_ERR_BASE boundary.

17 **Advice to users**

18 As previously stated, upon completing its work, and prior to returning, each handler *must* call the
19 event handler completion function provided when it was invoked (including a status code plus any
20 information to be passed to later handlers) so that the chain can continue being progressed. An
21 event handler can terminate all further progress along the chain by passing the
22 **PMIX_EVENT_ACTION_COMPLETE** status to the completion callback function. Note that the
23 parameters passed to the event handler (e.g., the *info* and *results* arrays) will cease to be valid once
24 the completion function has been called - thus, any information in the incoming parameters that
will be referenced following the call to the completion function must be copied.

25 **10.1.3 Event registration constants**

26 **PMIX_ERR_EVENT_REGISTRATION** -144 Error in event registration.

1 10.1.4 System events

2 **PMIX_EVENT_SYS_BASE** Mark the beginning of a dedicated range of constants for system
3 event reporting.

4 **PMIX_EVENT_NODE_DOWN** -231 A node has gone down - the identifier of the affected
5 node will be included in the notification.


6 **PMIX_EVENT_NODE_OFFLINE** -232 A node has been marked as *offline* - the identifier of
7 the affected node will be included in the notification.

8 **PMIX_EVENT_SYS_OTHER** Mark the end of a dedicated range of constants for system event
9 reporting.

10 Detect system event constant

11 Test a given event constant to see if it falls within the dedicated range of constants for system event
12 reporting.

PMIx v2.2



14 **IN** a

15 Error constant to be checked ([pmix_status_t](#))

16 Returns **true** if the provided values falls within the dedicated range of events for system event
17 reporting.

18 10.1.5 Event handler registration and notification attributes

19 Attributes to support event registration and notification.

20 **PMIX_EVENT_HDLR_NAME** "pmix.evname" (**char***)

21 String name identifying this handler.

22 **PMIX_EVENT_HDLR_FIRST** "pmix.evfirst" (**bool**)

23 Invoke this event handler before any other handlers.

24 **PMIX_EVENT_HDLR_LAST** "pmix.evlast" (**bool**)

25 Invoke this event handler after all other handlers have been called.

26 **PMIX_EVENT_HDLR_FIRST_IN_CATEGORY** "pmix.evfirstcat" (**bool**)

27 Invoke this event handler before any other handlers in this category.

28 **PMIX_EVENT_HDLR_LAST_IN_CATEGORY** "pmix.evlastcat" (**bool**)

29 Invoke this event handler after all other handlers in this category have been called.

30 **PMIX_EVENT_HDLR_BEFORE** "pmix.evbefore" (**char***)

31 Put this event handler immediately before the one specified in the (**char***) value.

32 **PMIX_EVENT_HDLR_AFTER** "pmix.evafter" (**char***)

33 Put this event handler immediately after the one specified in the (**char***) value.

34 **PMIX_EVENT_HDLR_PREPEND** "pmix.evprepend" (**bool**)

35 Prepend this handler to the precedence list within its category.

1 **PMIX_EVENT_HDLR_APPEND** "pmix.evappend" (bool)
 2 Append this handler to the precedence list within its category.

3 **PMIX_EVENT_CUSTOM_RANGE** "pmix.evrangle" (pmix_data_array_t*)
 4 Array of **pmix_proc_t** defining range of event notification.

5 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (pmix_proc_t)
 6 The single process that was affected.

7 **PMIX_EVENT_AFFECTED_PROCS** "pmix.evaffected" (pmix_data_array_t*)
 8 Array of **pmix_proc_t** defining affected processes.

9 **PMIX_EVENT_NON_DEFAULT** "pmix.evnondf" (bool)
 10 Event is not to be delivered to default event handlers.

11 **PMIX_EVENT_RETURN_OBJECT** "pmix.evobject" (void *)
 12 Object to be returned whenever the registered callback function **cbfunc** is invoked. The
 13 object will only be returned to the process that registered it.

14 **PMIX_EVENT_DO_NOT_CACHE** "pmix.evnocache" (bool)
 15 Instruct the PMIx server not to cache the event.

16 **PMIX_EVENT_PROXY** "pmix.evproxy" (pmix_proc_t*)
 17 PMIx server that sourced the event.

18 **PMIX_EVENT_TEXT_MESSAGE** "pmix.evtext" (char*)
 19 Text message suitable for output by recipient - e.g., describing the cause of the event.

20 **PMIX_EVENT_TIMESTAMP** "pmix.evtstamp" (time_t)
 21 System time when the associated event occurred.

22 10.1.5.1 Fault tolerance event attributes

23 The following attributes may be used by the host environment when providing an event notification
 24 as qualifiers indicating the action it intends to take in response to the event:

25 **PMIX_EVENT_TERMINATE_SESSION** "pmix.evterm.sess" (bool)
 26 The RM intends to terminate this session.

27 **PMIX_EVENT_TERMINATE_JOB** "pmix.evterm.job" (bool)
 28 The RM intends to terminate this job.

29 **PMIX_EVENT_TERMINATE_NODE** "pmix.evterm.node" (bool)
 30 The RM intends to terminate all processes on this node.

31 **PMIX_EVENT_TERMINATE_PROC** "pmix.evterm.proc" (bool)
 32 The RM intends to terminate just this process.

33 **PMIX_EVENT_ACTION_TIMEOUT** "pmix.evtimeout" (int)
 34 The time in seconds before the RM will execute the indicated operation.

35 10.1.5.2 Hybrid programming event attributes

36 The following attributes may be used by programming models to coordinate their use of common
 37 resources within a process in conjunction with the **PMIX_OPENMP_PARALLEL_ENTERED** event:

38 **PMIX_MODEL_PHASE_NAME** "pmix.mdl.phase" (char*)
 39 User-assigned name for a phase in the application execution (e.g., "cfd reduction").

40 **PMIX_MODEL_PHASE_TYPE** "pmix.mdl.ptype" (char*)
 41 Type of phase being executed (e.g., "matrix multiply").

1 10.1.6 Notification Function

2 Summary

3 The `pmix_notification_fn_t` is called by PMIx to deliver notification of an event.

Advice to users

4 The PMIx *ad hoc* v1.0 Standard defined an error notification function with an identical name, but
5 different signature than the v2.0 Standard described below. The *ad hoc* v1.0 version was removed
6 from the v2.0 Standard is not included in this document to avoid confusion.

PMIx v2.0

C

```
7 typedef void (*pmix_notification_fn_t)
8     (size_t evhdlr_registration_id,
9      pmix_status_t status,
10     const pmix_proc_t *source,
11     pmix_info_t info[], size_t ninfo,
12     pmix_info_t results[], size_t nresults,
13     pmix_event_notification_cbfunc_fn_t cbfunc,
14     void *cbdata);
```

C

15 **IN evhdlr_registration_id**
16 Registration number of the handler being called (`size_t`)

17 **IN status**
18 Status associated with the operation (`pmix_status_t`)

19 **IN source**
20 Identifier of the process that generated the event (`pmix_proc_t`). If the source is the SMS,
21 then the nspace will be empty and the rank will be `PMIX_RANK_UNDEF`

22 **IN info**
23 Information describing the event (`pmix_info_t`). This argument will be `NULL` if no
24 additional information was provided by the event generator.

25 **IN ninfo**
26 Number of elements in the info array (`size_t`)

27 **IN results**
28 Aggregated results from prior event handlers servicing this event (`pmix_info_t`). This
29 argument will be `NULL` if this is the first handler servicing the event, or if no prior handlers
30 provided results.

31 **IN nresults**
32 Number of elements in the results array (`size_t`)

33 **IN cbfunc**
34 `pmix_event_notification_cbfunc_fn_t` callback function to be executed upon
35 completion of the handler's operation and prior to handler return (function reference).

IN `cbdata`

Callback data to be passed to `cbfunc` (memory reference)

Description

Note that different RMs may provide differing levels of support for event notification to application processes. Thus, the *info* array may be **NULL** or may contain detailed information of the event. It is the responsibility of the application to parse any provided info array for defined key-values if it so desires.

Advice to users

Possible uses of the *info* array include:

- for the host RM to alert the process as to planned actions, such as aborting the session, in response to the reported event
- provide a timeout for alternative action to occur, such as for the application to request an alternate response to the event

For example, the RM might alert the application to the failure of a node that resulted in termination of several processes, and indicate that the overall session will be aborted unless the application requests an alternative behavior in the next 5 seconds. The application then has time to respond with a checkpoint request, or a request to recover from the failure by obtaining replacement nodes and restarting from some earlier checkpoint.

Support for these options is left to the discretion of the host RM. Info keys are included in the common definitions above but may be augmented by environment vendors.

Advice to PMIx server hosts

On the server side, the notification function is used to inform the PMIx server library's host of a detected event in the PMIx server library. Events generated by PMIx clients are communicated to the PMIx server library, but will be relayed to the host via the `pmix_server_notify_event_fn_t` function pointer, if provided.

10.1.7 `PMIx_Deregister_event_handler`

Summary

Deregister an event handler.

Format

C

```
pmix_status_t
PMIx_Deregister_event_handler(size_t evhdlr_ref,
                              pmix_op_cbfunc_t cbfunc,
                              void *cbdata);
```

C

IN `evhdlr_ref`

Event handler ID returned by registration (`size_t`)

IN `cbfunc`

Callback function to be executed upon completion of operation `pmix_op_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the `cbfunc` callback function (memory reference)

If `cbfunc` is `NULL`, the function will be treated as a *blocking* call and the result of the operation returned in the status code.

If `cbfunc` is non-`NULL`, the function will be treated as a *non-blocking* call.

A successful return indicates that the request is being processed and the result will be returned in the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning from the API. The callback function, `cbfunc`, is only called when `PMIX_SUCCESS` is returned.

- `PMIX_OPERATION_SUCCEEDED`, returned when the request was immediately processed successfully - the `cbfunc` will *not* be called.

The returned status code of `cbfunc` will be one of the following:

- `PMIX_SUCCESS` The event handler was successfully deregistered.
- `PMIX_ERR_BAD_PARAM` The provided `evhdlr_ref` was unrecognized.
- `PMIX_ERR_NOT_SUPPORTED` The PMIx implementation does not support event notification.

Description

Deregister an event handler. Note that no events corresponding to the referenced registration may be delivered following completion of the deregistration operation (either return from the API with `PMIX_OPERATION_SUCCEEDED` or execution of the `cbfunc`).

10.1.8 PMIx_Notify_event

Summary

Report an event for notification via any registered event handler.

Format

C

```
pmix_status_t
PMIx_Notify_event (pmix_status_t status,
                  const pmix_proc_t *source,
                  pmix_data_range_t range,
                  pmix_info_t info[], size_t ninfo,
                  pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

IN status

Status code of the event ([pmix_status_t](#))

IN source

Pointer to a [pmix_proc_t](#) identifying the original reporter of the event (handle)

IN range

Range across which this notification shall be delivered ([pmix_data_range_t](#))

IN info

Array of [pmix_info_t](#) structures containing any further info provided by the originator of the event (array of handles)

IN ninfo

Number of elements in the *info* array ([size_t](#))

IN cbfunc

Callback function to be executed upon completion of operation [pmix_op_cbfunc_t](#) (function reference)

IN cbdata

Data to be passed to the *cbfunc* callback function (memory reference)

If *cbfunc* is **NULL**, the function will be treated as a *blocking* call and the result of the operation returned in the status code.

If *cbfunc* is non-**NULL**, the function will be treated as a *non-blocking* call.

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned. Note that a successful call does *not* reflect the success or failure of delivering the event to any recipients.

Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- **PMIX_OPERATION_SUCCEEDED**, returned when the request was immediately processed successfully - the *cbfunc* will *not* be called.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_EVENT_NON_DEFAULT "pmix.evnondef" (bool)

Event is not to be delivered to default event handlers.

PMIX_EVENT_CUSTOM_RANGE "pmix.evrangle" (pmix_data_array_t*)

Array of pmix_proc_t defining range of event notification.

PMIX_EVENT_DO_NOT_CACHE "pmix.evnocache" (bool)

Instruct the PMIx server not to cache the event.

PMIX_EVENT_PROXY "pmix.evproxy" (pmix_proc_t*)

PMIx server that sourced the event.

PMIX_EVENT_TEXT_MESSAGE "pmix.evtext" (char*)

Text message suitable for output by recipient - e.g., describing the cause of the event.

Host environments that implement support for PMIx event notification are required to provide the following attributes for all events generated by the environment:

PMIX_EVENT_AFFECTED_PROC "pmix.evproc" (pmix_proc_t)

The single process that was affected.

PMIX_EVENT_AFFECTED_PROCS "pmix.evaffected" (pmix_data_array_t*)

Array of pmix_proc_t defining affected processes.

Optional Attributes

Host environments that support PMIx event notification may offer notifications for environmental events impacting the job and for SMS events relating to the job. The following attributes may optionally be included to indicate the host environment's intended response to the event:

PMIX_EVENT_TERMINATE_SESSION "pmix.evterm.sess" (bool)

The RM intends to terminate this session.

PMIX_EVENT_TERMINATE_JOB "pmix.evterm.job" (bool)

The RM intends to terminate this job.

PMIX_EVENT_TERMINATE_NODE "pmix.evterm.node" (bool)

The RM intends to terminate all processes on this node.

PMIX_EVENT_TERMINATE_PROC "pmix.evterm.proc" (bool)

The RM intends to terminate just this process.

PMIX_EVENT_ACTION_TIMEOUT "pmix.evtimeout" (int)

The time in seconds before the RM will execute the indicated operation.

▲-----▲

1 Description

2 Report an event for notification via any registered event handler. This function can be called by any
3 PMIx process, including application processes, PMIx servers, and SMS elements. The PMIx server
4 calls this API to report events it detected itself so that the host SMS daemon distribute and handle
5 them, and to pass events given to it by its host down to any attached client processes for processing.
6 Examples might include notification of the failure of another process, detection of an impending
7 node failure due to rising temperatures, or an intent to preempt the application. Events may be
8 locally generated or come from anywhere in the system.

9 Host SMS daemons call the API to pass events down to its embedded PMIx server both for
10 transmittal to local client processes and for the host's own internal processing where the host has
11 registered its own event handlers. The PMIx server library is not allowed to echo any event given to
12 it by its host via this API back to the host through the `pmix_server_notify_event_fn_t`
13 server module function. The host is required to deliver the event to all PMIx servers where the
14 targeted processes either are currently running, or (if they haven't started yet) might be running at
15 some point in the future as the events are required to be cached by the PMIx server library.

16 Client application processes can call this function to notify the SMS and/or other application
17 processes of an event it encountered. Note that processes are not constrained to report status values
18 defined in the official PMIx standard — any integer value can be used. Thus, applications are free
19 to define their own internal events and use the notification system for their own internal purposes.

▼-----▼

Advice to users

20 The callback function will be called upon completion of the `notify_event` function's actions.
21 At that time, any messages required for executing the operation (e.g., to send the notification to the
22 local PMIx server) will have been queued, but may not yet have been transmitted. The caller is
23 required to maintain the input data until the callback function has been executed — the sole purpose
24 of the callback function is to indicate when the input data is no longer required.

▲-----▲

1 10.1.9 Notification Handler Completion Callback Function

2 Summary

3 The `pmix_event_notification_cbfunc_fn_t` is called by event handlers to indicate
4 completion of their operations.

```
5 typedef void (*pmix_event_notification_cbfunc_fn_t)  
6     (pmix_status_t status,  
7      pmix_info_t *results, size_t nresults,  
8      pmix_op_cbfunc_t cbfunc, void *thiscbdata,  
9      void *notification_cbdata);
```

10 **IN status**
11 Status returned by the event handler's operation (`pmix_status_t`)

12 **IN results**
13 Results from this event handler's operation on the event (`pmix_info_t`)

14 **IN nresults**
15 Number of elements in the results array (`size_t`)

16 **IN cbfunc**
17 `pmix_op_cbfunc_t` function to be executed when PMIx completes processing the
18 callback (function reference)

19 **IN thiscbdata**
20 Callback data that was passed in to the handler (memory reference)

21 **IN cbdata**
22 Callback data to be returned when PMIx executes cbfunc (memory reference)

23 Description

24 Define a callback by which an event handler can notify the PMIx library that it has completed its
25 response to the notification. The handler is *required* to execute this callback so the library can
26 determine if additional handlers need to be called. The handler shall return
27 `PMIX_EVENT_ACTION_COMPLETE` if no further action is required. The return status of each
28 event handler and any returned `pmix_info_t` structures will be added to the *results* array of
29 `pmix_info_t` passed to any subsequent event handlers to help guide their operation.

30 If non-NULL, the provided callback function will be called to allow the event handler to release the
31 provided info array and execute any other required cleanup operations.

32 10.1.9.1 Completion Callback Function Status Codes

33 The following status code may be returned indicating various actions taken by other event handlers.

34 `PMIX_EVENT_NO_ACTION_TAKEN` -331 Event handler: No action taken.
35 `PMIX_EVENT_PARTIAL_ACTION_TAKEN` -332 Event handler: Partial action taken.
36 `PMIX_EVENT_ACTION_DEFERRED` -333 Event handler: Action deferred.
37 `PMIX_EVENT_ACTION_COMPLETE` -334 Event handler: Action complete.

CHAPTER 11

Data Packing and Unpacking

1 PMIx intentionally does not include support for internode communications in the standard, instead
2 relying on its host SMS environment to transfer any needed data and/or requests between nodes.
3 These operations frequently involve PMIx-defined public data structures that include binary data.
4 Many HPC clusters are homogeneous, and so transferring the structures can be done rather simply.
5 However, greater effort is required in heterogeneous environments to ensure binary data is correctly
6 transferred. PMIx buffer manipulation functions are provided for this purpose via standardized
7 interfaces to ease adoption.

8 11.1 Data Buffer Type

9 The `pmix_data_buffer_t` structure describes a data buffer used for packing and unpacking.

PMIx v2.0

C

```
10 typedef struct pmix_data_buffer {  
11     /** Start of my memory */  
12     char *base_ptr;  
13     /** Where the next data will be packed to  
14         (within the allocated memory starting  
15         at base_ptr) */  
16     char *pack_ptr;  
17     /** Where the next data will be unpacked  
18         from (within the allocated memory  
19         starting as base_ptr) */  
20     char *unpack_ptr;  
21     /** Number of bytes allocated (starting  
22         at base_ptr) */  
23     size_t bytes_allocated;  
24     /** Number of bytes used by the buffer  
25         (i.e., amount of data - including  
26         overhead - packed in the buffer) */  
27     size_t bytes_used;  
28 } pmix_data_buffer_t;
```

C

1 11.2 Support Macros

2 PMIx provides a set of convenience macros for creating, initiating, and releasing data buffers.

3 **Static initializer for the data buffer structure**

4 *Provisional* Provide a static initializer for the `pmix_data_buffer_t` fields.

PMIx v4.2

▼ C _____ ▼

5 **PMIX_DATA_BUFFER_STATIC_INIT**

▲ C _____ ▲

6 **PMIX_DATA_BUFFER_CREATE**

7 Allocate memory for a `pmix_data_buffer_t` object and initialize it. This macro uses `calloc` to
8 allocate memory for the buffer and initialize all fields in it

PMIx v2.0

▼ C _____ ▼

9 **PMIX_DATA_BUFFER_CREATE(buffer);**

▲ C _____ ▲

10 **OUT buffer**

11 Variable to be assigned the pointer to the allocated `pmix_data_buffer_t` (handle)

12 **PMIX_DATA_BUFFER_RELEASE**

13 Free a `pmix_data_buffer_t` object and the data it contains. Free's the data contained in the
14 buffer, and then free's the buffer itself

PMIx v2.0

▼ C _____ ▼

15 **PMIX_DATA_BUFFER_RELEASE(buffer);**

▲ C _____ ▲

16 **IN buffer**

17 Pointer to the `pmix_data_buffer_t` to be released (handle)

18 **PMIX_DATA_BUFFER_CONSTRUCT**

19 Initialize a statically declared `pmix_data_buffer_t` object.

PMIx v2.0

▼ C _____ ▼

20 **PMIX_DATA_BUFFER_CONSTRUCT(buffer);**

▲ C _____ ▲

21 **IN buffer**

22 Pointer to the allocated `pmix_data_buffer_t` that is to be initialized (handle)

1 **PMIX_DATA_BUFFER_DESTRUCT**
2 Release the data contained in a `pmix_data_buffer_t` object.

▼ `PMIX_DATA_BUFFER_DESTRUCT` C `pmix_data_buffer_t` ▼
3 **PMIX_DATA_BUFFER_DESTRUCT**(buffer);
▲ `PMIX_DATA_BUFFER_DESTRUCT` C `pmix_data_buffer_t` ▲

4 **IN** `buffer`
5 Pointer to the `pmix_data_buffer_t` whose data is to be released (handle)

6 **PMIX_DATA_BUFFER_LOAD**
7 Load a blob into a `pmix_data_buffer_t` object. Load the given data into the provided
8 `pmix_data_buffer_t` object, usually done in preparation for unpacking the provided data.
9 Note that the data is *not* copied into the buffer - thus, the blob must not be released until after
10 operations on the buffer have completed.

PMIx v2.0 ▼ `PMIX_DATA_BUFFER_LOAD` C `pmix_data_buffer_t` ▼
11 **PMIX_DATA_BUFFER_LOAD**(buffer, data, size);
▲ `PMIX_DATA_BUFFER_LOAD` C `pmix_data_buffer_t` ▲

12 **IN** `buffer`
13 Pointer to a pre-allocated `pmix_data_buffer_t` (handle)

14 **IN** `data`
15 Pointer to a blob (`char*`)

16 **IN** `size`
17 Number of bytes in the blob `size_t`

18 **PMIX_DATA_BUFFER_UNLOAD**
19 Unload the data from a `pmix_data_buffer_t` object. Extract the data in a buffer, assigning the
20 pointer to the data (and the number of bytes in the blob) to the provided variables, usually done to
21 transmit the blob to a remote process for unpacking. The buffer's internal pointer will be set to
22 NULL to protect the data upon buffer destruct or release - thus, the user is responsible for releasing
23 the blob when done with it.

PMIx v2.0 ▼ `PMIX_DATA_BUFFER_UNLOAD` C `pmix_data_buffer_t` ▼
24 **PMIX_DATA_BUFFER_UNLOAD**(buffer, data, size);
▲ `PMIX_DATA_BUFFER_UNLOAD` C `pmix_data_buffer_t` ▲

25 **IN** `buffer`
26 Pointer to the `pmix_data_buffer_t` whose data is to be extracted (handle)

27 **OUT** `data`
28 Variable to be assigned the pointer to the extracted blob (`void*`)

29 **OUT** `size`
30 Variable to be assigned the number of bytes in the blob `size_t`

1 11.3 General Routines

2 The following routines are provided to support internode transfers in heterogeneous environments.

3 11.3.1 PMIx_Data_pack

4 Summary

5 Pack one or more values of a specified type into a buffer, usually for transmission to another process.

6 *PMIx v2.0* Format

C

```
7 pmix_status_t
8 PMIx_Data_pack(const pmix_proc_t *target,
9               pmix_data_buffer_t *buffer,
10              void *src, int32_t num_vals,
11              pmix_data_type_t type);
```

C

12 IN target

13 Pointer to a [pmix_proc_t](#) containing the nspace/rank of the process that will be unpacking
14 the final buffer. A NULL value may be used to indicate that the target is based on the same
15 PMIx version as the caller. Note that only the target's nspace is relevant. (handle)

16 IN buffer

17 Pointer to a [pmix_data_buffer_t](#) where the packed data is to be stored (handle)

18 IN src

19 Pointer to a location where the data resides. Strings are to be passed as (char **) — i.e., the
20 caller must pass the address of the pointer to the string as the (void*). This allows the caller to
21 pass multiple strings in a single call. (memory reference)

22 IN num_vals

23 Number of elements pointed to by the *src* pointer. A string value is counted as a single value
24 regardless of length. The values must be contiguous in memory. Arrays of pointers (e.g.,
25 string arrays) should be contiguous, although the data pointed to need not be contiguous
26 across array entries. ([int32_t](#))

27 IN type

28 The type of the data to be packed ([pmix_data_type_t](#))

29 Returns PMIX_SUCCESS or one of the following error codes when the condition described occurs:

30 [PMIX_ERR_BAD_PARAM](#) The provided buffer or src is **NULL**

31 [PMIX_ERR_UNKNOWN_DATA_TYPE](#) The specified data type is not known to this
32 implementation

33 [PMIX_ERR_OUT_OF_RESOURCE](#) Not enough memory to support the operation

34 If none of the above return codes are appropriate, then an implementation must return either a
35 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Description

The pack function packs one or more values of a specified type into the specified buffer. The buffer must have already been initialized via the `PMIX_DATA_BUFFER_CREATE` or `PMIX_DATA_BUFFER_CONSTRUCT` macros — otherwise, `PMIx_Data_pack` will return an error. Providing an unsupported type flag will likewise be reported as an error.

Note that any data to be packed that is not hard type cast (i.e., not type cast to a specific size) may lose precision when unpacked by a non-homogeneous recipient. The `PMIx_Data_pack` function will do its best to deal with heterogeneity issues between the packer and unpacker in such cases. Sending a number larger than can be handled by the recipient will return an error code (generated upon unpacking) — the error cannot be detected during packing.

The namespace of the intended recipient of the packed buffer (i.e., the process that will be unpacking it) is used solely to resolve any data type differences between PMIx versions. The recipient must, therefore, be known to the user prior to calling the pack function so that the PMIx library is aware of the version the recipient is using. Note that all processes in a given namespace are *required* to use the same PMIx version — thus, the caller must only know at least one process from the target's namespace.

11.3.2 PMIx_Data_unpack

Summary

Unpack values from a `pmix_data_buffer_t`

Format

PMIx v2.0

`pmix_status_t`

```
PMIx_Data_unpack(const pmix_proc_t *source,  
                 pmix_data_buffer_t *buffer, void *dest,  
                 int32_t *max_num_values,  
                 pmix_data_type_t type);
```

IN source

Pointer to a `pmix_proc_t` structure containing the nspace/rank of the process that packed the provided buffer. A NULL value may be used to indicate that the source is based on the same PMIx version as the caller. Note that only the source's nspace is relevant. (handle)

IN buffer

A pointer to the buffer from which the value will be extracted. (handle)

INOUT dest

A pointer to the memory location into which the data is to be stored. Note that these values will be stored contiguously in memory. For strings, this pointer must be to (`char**`) to provide a means of supporting multiple string operations. The unpack function will allocate memory for each string in the array - the caller must only provide adequate memory for the array of pointers. (`void*`)

1 **INOUT** `max_num_values`

2 The number of values to be unpacked — upon completion, the parameter will be set to the
3 actual number of values unpacked. In most cases, this should match the maximum number
4 provided in the parameters — but in no case will it exceed the value of this parameter. Note
5 that unpacking fewer values than are actually available will leave the buffer in an unpackable
6 state — the function will return an error code to warn of this condition. (`int32_t`)

7 **IN** `type`

8 The type of the data to be unpacked — must be one of the PMIx defined data types
9 (`pmix_data_type_t`)

10 Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

11 `PMIX_ERR_BAD_PARAM` The provided buffer or dest is `NULL`

12 `PMIX_ERR_UNKNOWN_DATA_TYPE` The specified data type is not known to this
13 implementation

14 `PMIX_ERR_OUT_OF_RESOURCE` Not enough memory to support the operation

15 If none of the above return codes are appropriate, then an implementation must return either a
16 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

17 **Description**

18 The `unpack` function unpacks the next value (or values) of a specified type from the given buffer.
19 The buffer must have already been initialized via an `PMIX_DATA_BUFFER_CREATE` or
20 `PMIX_DATA_BUFFER_CONSTRUCT` call (and assumedly filled with some data) — otherwise, the
21 `unpack_value` function will return an error. Providing an unsupported type flag will likewise be
22 reported as an error, as will specifying a data type that *does not* match the type of the next item in
23 the buffer. An attempt to read beyond the end of the stored data held in the buffer will also return an
24 error.

25 Note that it is possible for the buffer to be corrupted and that PMIx will *think* there is a proper
26 variable type at the beginning of an unpack region — but that the value is bogus (e.g., just a byte
27 field in a string array that so happens to have a value that matches the specified data type flag).
28 Therefore, the data type error check is *not* completely safe.

29 Unpacking values is a "nondestructive" process — i.e., the values are not removed from the buffer.
30 It is therefore possible for the caller to re-unpack a value from the same buffer by resetting the
31 `unpack_ptr`.

32 **Warning:** The caller is responsible for providing adequate memory storage for the requested data.
33 The user must provide a parameter indicating the maximum number of values that can be unpacked
34 into the allocated memory. If more values exist in the buffer than can fit into the memory storage,
35 then the function will unpack what it can fit into that location and return an error code indicating
36 that the buffer was only partially unpacked.

37 Note that any data that was not hard type cast (i.e., not type cast to a specific size) when packed may
38 lose precision when unpacked by a non-homogeneous recipient. PMIx will do its best to deal with
39 heterogeneity issues between the packer and unpacker in such cases. Sending a number larger than

1 can be handled by the recipient will return an error code generated upon unpacking — these errors
2 cannot be detected during packing.

3 The namespace of the process that packed the buffer is used solely to resolve any data type
4 differences between PMIx versions. The packer must, therefore, be known to the user prior to
5 calling the pack function so that the PMIx library is aware of the version the packer is using. Note
6 that all processes in a given namespace are *required* to use the same PMIx version — thus, the
7 caller must only know at least one process from the packer’s namespace.

8 11.3.3 PMIx_Data_copy

9 Summary

10 Copy a data value from one location to another.

11 Format

PMIx v2.0

C

```
12 pmix_status_t  
13 PMIx_Data_copy(void **dest, void *src,  
14                 pmix_data_type_t type);
```

C

15 IN dest

16 The address of a pointer into which the address of the resulting data is to be stored. (**void****)

17 IN src

18 A pointer to the memory location from which the data is to be copied (handle)

19 IN type

20 The type of the data to be copied — must be one of the PMIx defined data types.
21 (**pmix_data_type_t**)

22 Returns PMIX_SUCCESS or one of the following error codes when the condition described occurs:

23 **PMIX_ERR_BAD_PARAM** The provided src or dest is **NULL**

24 **PMIX_ERR_UNKNOWN_DATA_TYPE** The specified data type is not known to this
25 implementation

26 **PMIX_ERR_OUT_OF_RESOURCE** Not enough memory to support the operation

27 If none of the above return codes are appropriate, then an implementation must return either a
28 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

29 Description

30 Since registered data types can be complex structures, the system needs some way to know how to
31 copy the data from one location to another (e.g., for storage in the registry). This function, which
32 can call other copy functions to build up complex data types, defines the method for making a copy
33 of the specified data type.

1 11.3.4 PMIx_Data_print

2 Summary

3 Pretty-print a data value.

4 Format

PMIx v2.0

```
pmix_status_t  
PMIx_Data_print(char **output, char *prefix,  
                void *src, pmix_data_type_t type);
```

8 IN output

9 The address of a pointer into which the address of the resulting output is to be stored.

10 (**char****)

11 IN prefix

12 String to be prepended to the resulting output (**char***)

13 IN src

14 A pointer to the memory location of the data value to be printed (handle)

15 IN type

16 The type of the data value to be printed — must be one of the PMIx defined data types.

17 ([pmix_data_type_t](#))

18 Returns PMIX_SUCCESS or one of the following error codes when the condition described occurs:

19 [PMIX_ERR_BAD_PARAM](#) The provided data type is not recognized.

20 If none of the above return codes are appropriate, then an implementation must return either a
21 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

22 Description

23 Since registered data types can be complex structures, the system needs some way to know how to
24 print them (i.e., convert them to a string representation). Primarily for debug purposes.

25 11.3.5 PMIx_Data_copy_payload

26 Summary

27 Copy a payload from one buffer to another

Format

C

`pmix_status_t`

```
PMIx_Data_copy_payload(pmix_data_buffer_t *dest,  
                      pmix_data_buffer_t *src);
```

C

IN `dest`

Pointer to the destination `pmix_data_buffer_t` (handle)

IN `src`

Pointer to the source `pmix_data_buffer_t` (handle)

Returns one of the following:

PMIX_SUCCESS The data has been copied as requested

PMIX_ERR_BAD_PARAM The `src` and `dest` `pmix_data_buffer_t` types do not match

PMIX_ERR_NOT_SUPPORTED The PMIx implementation does not support this function.

Description

This function will append a copy of the payload in one buffer into another buffer. Note that this is *not* a destructive procedure — the source buffer’s payload will remain intact, as will any pre-existing payload in the destination’s buffer. Only the unpacked portion of the source payload will be copied.

11.3.6 PMIx_Data_load

Summary

Load a buffer with the provided payload

Format

C

`pmix_status_t`

```
PMIx_Data_load(pmix_data_buffer_t *dest,  
              pmix_byte_object_t *src);
```

C

IN `dest`

Pointer to the destination `pmix_data_buffer_t` (handle)

IN `src`

Pointer to the source `pmix_byte_object_t` (handle)

Returns one of the following:

PMIX_SUCCESS The data has been loaded as requested

PMIX_ERR_BAD_PARAM The `dest` structure pointer is **NULL**

PMIX_ERR_NOT_SUPPORTED The PMIx implementation does not support this function.

Provisional
v4.1

Description

The load function allows the caller to transfer the contents of the *src* `pmix_byte_object_t` to the *dest* target buffer. If a payload already exists in the buffer, the function will "free" the existing data to release it, and then replace the data payload with the one provided by the caller.

Advice to users

The buffer must be allocated or constructed in advance - failing to do so will cause the load function to return an error code.

The caller is responsible for pre-packing the provided payload. For example, the load function cannot convert to network byte order any data contained in the provided payload.

11.3.7 PMIx_Data_unload

Summary

Unload a buffer into a byte object

Format

Provisional
v4.1

```
pmix_status_t  
PMIx_Data_unload(pmix_data_buffer_t *src,  
                 pmix_byte_object_t *dest);
```

IN `src`

Pointer to the source `pmix_data_buffer_t` (handle)

IN `dest`

Pointer to the destination `pmix_byte_object_t` (handle)

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

PMIX_ERR_BAD_PARAM The destination and/or source pointer is **NULL**

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Description

The unload function provides the caller with a pointer to the portion of the data payload within the buffer that has not yet been unpacked, along with the size of that region. Any portion of the payload that was previously unpacked using the `PMIx_Data_unpack` routine will be ignored. This allows the user to directly access the payload.

Advice to users

This is a destructive operation. While the payload returned in the destination `pmix_byte_object_t` is undisturbed, the function will clear the `src`'s pointers to the payload. Thus, the `src` and the payload are completely separated, leaving the caller able to free or destruct the `src`.

11.3.8 PMIx_Data_compress

Summary

Perform a lossless compression on the provided data

Format

C

```
bool
PMIx_Data_compress(const uint8_t *inbytes, size_t size,
                  uint8_t **outbytes, size_t *nbytes);
```

C

IN inbytes

Pointer to the source data (handle)

IN size

Number of bytes in the source data region (`size_t`)

OUT outbytes

Address where the pointer to the compressed data region is to be returned (handle)

OUT nbytes

Address where the number of bytes in the compressed data region is to be returned (handle)

Returns one of the following:

- **True** The data has been compressed as requested
- **False** The data has not been compressed

Provisional
v4.1

Description

Compress the provided data block. Destination memory will be allocated if operation is successfully concluded. Caller is responsible for release of the allocated region. The input data block will remain unaltered.

Note: the compress function will return **False** if the operation would not result in a smaller data block.

11.3.9 PMIx_Data_decompress

Summary

Decompress the provided data

Format

Provisional
v4.1

C

bool

```
PMIx_Data_decompress(const uint8_t *inbytes, size_t size,  
                    uint8_t **outbytes, size_t *nbytes);
```

C

OUT outbytes

Address where the pointer to the decompressed data region is to be returned (handle)

OUT nbytes

Address where the number of bytes in the decompressed data region is to be returned (handle)

IN inbytes

Pointer to the source data (handle)

IN size

Number of bytes in the source data region (**size_t**)

Returns one of the following:

- **True** The data has been decompressed as requested
- **False** The data has not been decompressed

Description

Decompress the provided data block. Destination memory will be allocated if operation is successfully concluded. Caller is responsible for release of the allocated region. The input data block will remain unaltered.

Only data compressed by the **PMIx_Data_compress** API can be decompressed by this function. Passing data that has not been compressed by **PMIx_Data_compress** will lead to unexpected and potentially catastrophic results.

1 11.3.10 `PMIx_Data_embed`

2 Summary

3 Embed a data payload into a buffer

4 Format

```
5 pmix_status_t  
6 PMIx_Data_embed(pmix_data_buffer_t *buffer,  
7                 const pmix_byte_object_t *payload);
```

8 OUT `buffer`

9 Address of the buffer where the payload is to be embedded (handle)

10 IN `payload`

11 Address of the `pmix_byte_object_t` structure containing the data to be embedded into
12 the buffer (handle)

13 Returns one of the following:

14 **PMIX_SUCCESS** The data has been embedded as requested

15 **PMIX_ERR_BAD_PARAM** The destination and/or source pointer is **NULL**

16 **PMIX_ERR_NOT_SUPPORTED** The PMIx implementation does not support this function.

17 Description

18 The embed function is identical in operation to `PMIx_Data_load` except that it does *not* clear
19 the payload object upon completion.

CHAPTER 12

Process Management

1 This chapter defines functionality processes can use to abort processes, spawn processes, and
2 determine the relative locality of local processes.

3 12.1 Abort

4 PMIx provides a dedicated API by which an application can request that specified processes be
5 aborted by the system.

6 12.1.1 PMIx_Abort

7 Summary

8 Abort the specified processes

9 *PMIx v1.0* Format

```
10 pmix_status_t  
11 PMIx_Abort(int status, const char msg[],  
12           pmix_proc_t procs[], size_t nprocs)
```

13 IN status

14 Error code to return to invoking environment (integer)

15 IN msg

16 String message to be returned to user (string)

17 IN procs

18 Array of `pmix_proc_t` structures (array of handles)

19 IN nprocs

20 Number of elements in the *procs* array (integer)

21 A successful return indicates that the requested processes are in a terminated state. Note that the
22 function shall not return in this situation if the caller's own process was included in the request.

23 Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- **PMIX_ERR_PARAM_VALUE_NOT_SUPPORTED** if the PMIx implementation and host environment support this API, but the request includes processes that the host environment cannot abort - e.g., if the request is to abort subsets of processes from a namespace, or processes outside of the caller's own namespace, and the host environment does not permit such operations. In this case, none of the specified processes will be terminated.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Description

Request that the host resource manager print the provided message and abort the provided array of *procs*. A Unix or POSIX environment should handle the provided status as a return error code from the main program that launched the application. A **NULL** for the *procs* array indicates that all processes in the caller's namespace are to be aborted, including itself - this is the equivalent of passing a **pmix_proc_t** array element containing the caller's namespace and a rank value of **PMIX_RANK_WILDCARD**. While it is permitted for a caller to request abort of processes from namespaces other than its own, not all environments will support such requests. Passing a **NULL** *msg* parameter is allowed.

The function shall not return until the host environment has carried out the operation on the specified processes. If the caller is included in the array of targets, then the function will not return unless the host is unable to execute the operation.

Advice to users

The response to this request is somewhat dependent on the specific RM and its configuration (e.g., some resource managers will not abort the application if the provided status is zero unless specifically configured to do so, some cannot abort subsets of processes in an application, and some may not permit termination of processes outside of the caller's own namespace), and thus lies outside the control of PMIx itself. However, the PMIx client library shall inform the RM of the request that the specified *procs* be aborted, regardless of the value of the provided status.

Note that race conditions caused by multiple processes calling **PMIx_Abort** are left to the server implementation to resolve with regard to which status is returned and what messages (if any) are printed.

12.2 Process Creation

The **PMIx_Spawn** commands spawn new processes and/or applications in the PMIx universe. This may include requests to extend the existing resource allocation or obtain a new one, depending upon provided and supported attributes.

1 12.2.1 PMIx_Spawn

2 Summary

3 Spawn a new job.

4 Format

PMIx v1.0

C

```
5 pmix_status_t  
6 PMIx_Spawn(const pmix_info_t job_info[], size_t ninfo,  
7             const pmix_app_t apps[], size_t napps,  
8             char nspace[])
```

C

9 **IN** **job_info**
10 Array of info structures (array of handles)

11 **IN** **ninfo**
12 Number of elements in the *job_info* array (integer)

13 **IN** **apps**
14 Array of **pmix_app_t** structures (array of handles)

15 **IN** **napps**
16 Number of elements in the *apps* array (integer)

17 **OUT** **nspace**
18 Namespace of the new job (string)

19 Returns **PMIX_SUCCESS** or a negative value indicating the error.

Required Attributes

20 PMIx libraries are not required to directly support any attributes for this function. However, any
21 provided attributes must be passed to the host environment for processing.

22 Host environments are required to support the following attributes when present in either the
23 *job_info* or the *info* array of an element of the *apps* array:

24 **PMIX_WDIR** "**pmix.wdir**" (**char***)
25 Working directory for spawned processes.

26 **PMIX_SET_SESSION_CWD** "**pmix.ssn cwd**" (**bool**)
27 Set the current working directory to the session working directory assigned by the RM - can
28 be assigned to the entire job (by including attribute in the *job_info* array) or on a
29 per-application basis in the *info* array for each **pmix_app_t**.

30 **PMIX_PREFIX** "**pmix.prefix**" (**char***)
31 Prefix to use for starting spawned processes - i.e., the directory where the executables can be
32 found.

33 **PMIX_HOST** "**pmix.host**" (**char***)
34 Comma-delimited list of hosts to use for spawned processes.

1 **PMIX_HOSTFILE** "pmix.hostfile" (char*)

2 Hostfile to use for spawned processes.

3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_ADD_HOSTFILE "pmix.addhostfile" (char*)

Hostfile containing hosts to add to existing allocation.

PMIX_ADD_HOST "pmix.addhost" (char*)

Comma-delimited list of hosts to add to the allocation.

PMIX_PRELOAD_BIN "pmix.preloadbin" (bool)

Preload executables onto nodes prior to executing launch procedure.

PMIX_PRELOAD_FILES "pmix.preloadfiles" (char*)

Comma-delimited list of files to pre-position on nodes prior to executing launch procedure.

PMIX_PERSONALITY "pmix.pers" (char*)

Name of personality corresponding to programming model used by application - supported values depend upon PMIx implementation.

PMIX_DISPLAY_MAP "pmix.dispmap" (bool)

Display process mapping upon spawn.

PMIX_PPR "pmix.ppr" (char*)

Number of processes to spawn on each identified resource.

PMIX_MAPBY "pmix.mapby" (char*)

Process mapping policy - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the provided namespace. Supported values are launcher specific.

PMIX_RANKBY "pmix.rankby" (char*)

Process ranking policy - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the provided namespace. Supported values are launcher specific.

PMIX_BINDTO "pmix.bindto" (char*)

Process binding policy - when accessed using **PMIx_Get**, use the **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the provided namespace. Supported values are launcher specific.

PMIX_STDIN_TGT "pmix.stdin" (uint32_t)

Spawned process rank that is to receive any forwarded **stdin**.

PMIX_TAG_OUTPUT "pmix.tagout" (bool)

1 Tag **stdout**/**stderr** with the identity of the source process - can be assigned to the entire
2 job (by including attribute in the *job_info* array) or on a per-application basis in the *info*
3 array for each **pmix_app_t**.

4 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)

5 Timestamp output - can be assigned to the entire job (by including attribute in the *job_info*
6 array) or on a per-application basis in the *info* array for each **pmix_app_t**.

7 **PMIX_MERGE_STDERR_STDOUT** "pmix.mergeerrout" (bool)

8 Merge **stdout** and **stderr** streams - can be assigned to the entire job (by including
9 attribute in the *job_info* array) or on a per-application basis in the *info* array for each
10 **pmix_app_t**.

11 **PMIX_OUTPUT_TO_FILE** "pmix.outfile" (char*)

12 Direct output (both **stdout** and **stderr**) into files of form "<filename>.rank" - can be
13 assigned to the entire job (by including attribute in the *job_info* array) or on a per-application
14 basis in the *info* array for each **pmix_app_t**.

15 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)

16 Mark the **argv** with the rank of the process.

17 **PMIX_CPUS_PER_PROC** "pmix.cpusperproc" (uint32_t)

18 Number of PUs to assign to each rank - when accessed using **PMIx_Get**, use the
19 **PMIX_RANK_WILDCARD** value for the rank to discover the PUs/process assigned to the
20 provided namespace.

21 **PMIX_NO_PROCS_ON_HEAD** "pmix.nolocal" (bool)

22 Do not place processes on the head node.

23 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)

24 Do not oversubscribe the nodes - i.e., do not place more processes than allocated slots on a
25 node.

26 **PMIX_REPORT_BINDINGS** "pmix.repbind" (bool)

27 Report bindings of the individual processes.

28 **PMIX_CPU_LIST** "pmix.cpulist" (char*)

29 List of PUs to use for this job - when accessed using **PMIx_Get**, use the
30 **PMIX_RANK_WILDCARD** value for the rank to discover the PU list used for the provided
31 namespace.

32 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)

33 Application supports recoverable operations.

34 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)

35 Application is continuous, all failed processes should be immediately restarted.

36 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)

1 Maximum number of times to restart a process - when accessed using `PMIx_Get`, use the
2 `PMIX_RANK_WILDCARD` value for the rank to discover the max restarts for the provided
3 namespace.

4 `PMIX_SET_ENVAR "pmix.envar.set" (pmix_envar_t*)`
5 Set the envar to the given value, overwriting any pre-existing one

6 `PMIX_UNSET_ENVAR "pmix.envar.unset" (char*)`
7 Unset the environment variable specified in the string.

8 `PMIX_ADD_ENVAR "pmix.envar.add" (pmix_envar_t*)`
9 Add the environment variable, but do not overwrite any pre-existing one

10 `PMIX_PREPEND_ENVAR "pmix.envar.prepend" (pmix_envar_t*)`
11 Prepend the given value to the specified environmental value using the given separator
12 character, creating the variable if it doesn't already exist

13 `PMIX_APPEND_ENVAR "pmix.envar.append" (pmix_envar_t*)`
14 Append the given value to the specified environmental value using the given separator
15 character, creating the variable if it doesn't already exist

16 `PMIX_FIRST_ENVAR "pmix.envar.first" (pmix_envar_t*)`
17 Ensure the given value appears first in the specified envar using the separator character,
18 creating the envar if it doesn't already exist

19 `PMIX_ALLOC_QUEUE "pmix.alloc.queue" (char*)`
20 Name of the WLM queue to which the allocation request is to be directed, or the queue being
21 referenced in a query.

22 `PMIX_ALLOC_TIME "pmix.alloc.time" (uint32_t)`
23 Total session time (in seconds) being requested in an allocation request.

24 `PMIX_ALLOC_NUM_NODES "pmix.alloc.nnodes" (uint64_t)`
25 The number of nodes being requested in an allocation request.

26 `PMIX_ALLOC_NODE_LIST "pmix.alloc.nlist" (char*)`
27 Regular expression of the specific nodes being requested in an allocation request.

28 `PMIX_ALLOC_NUM_CPUS "pmix.alloc.ncpus" (uint64_t)`
29 Number of PUs being requested in an allocation request.

30 `PMIX_ALLOC_NUM_CPU_LIST "pmix.alloc.ncpulist" (char*)`
31 Regular expression of the number of PUs for each node being requested in an allocation
32 request.

33 `PMIX_ALLOC_CPU_LIST "pmix.alloc.cpulist" (char*)`
34 Regular expression of the specific PUs being requested in an allocation request.

35 `PMIX_ALLOC_MEM_SIZE "pmix.alloc.msize" (float)`

1 Number of Megabytes[base2] of memory (per process) being requested in an allocation
2 request.

3 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)

4 Fabric bandwidth (in Megabits[base2]/sec) for the job being requested in an allocation
5 request.

6 **PMIX_ALLOC_FABRIC_QOS** "pmix.alloc.netqos" (char*)

7 Fabric quality of service level for the job being requested in an allocation request.

8 **PMIX_ALLOC_FABRIC_TYPE** "pmix.alloc.nettype" (char*)

9 Type of desired transport (e.g., "tcp", "udp") being requested in an allocation request.

10 **PMIX_ALLOC_FABRIC_PLANE** "pmix.alloc.netplane" (char*)

11 ID string for the *fabric plane* to be used for the requested allocation.

12 **PMIX_ALLOC_FABRIC_ENDPTS** "pmix.alloc.endpts" (size_t)

13 Number of endpoints to allocate per *process* in the job.

14 **PMIX_ALLOC_FABRIC_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)

15 Number of endpoints to allocate per *node* for the job.

16 **PMIX_COSPAWN_APP** "pmix.cospawn" (bool)

17 Designated application is to be spawned as a disconnected job - i.e., the launcher shall not
18 include the application in any of the job-level values (e.g., **PMIX_RANK** within the job)
19 provided to any other application process generated by the same spawn request. Typically
20 used to cospawn debugger daemons alongside an application.

21 **PMIX_SPAWN_TOOL** "pmix.spwn.tool" (bool)

22 Indicate that the job being spawned is a tool.

23 **PMIX_EVENT_SILENT_TERMINATION** "pmix.avsilentterm" (bool)

24 Do not generate an event when this job normally terminates.

25 **PMIX_ENVARS_HARVESTED** "pmix.evar.hvstd" (bool)

26 Environmental parameters have been harvested by the spawn requestor - the server does not
27 need to harvest them.

28 **PMIX_JOB_TIMEOUT** "pmix.job.time" (int)

29 Time in seconds before the spawned job should time out and be terminated (0 => infinite),
30 defined as the total runtime of the job (equivalent to the walltime limit of typical batch
31 schedulers).

32 **PMIX_SPAWN_TIMEOUT** "pmix.sp.time" (int)

33 Time in seconds before spawn operation should time out (0 => infinite). Logically
34 equivalent to passing the **PMIX_TIMEOUT** attribute to the **PMIx_Spawn** API, it is
35 provided as a separate attribute to distinguish it from the **PMIX_JOB_TIMEOUT** attribute

Description

Spawn a new job. The assigned namespace of the spawned applications is returned in the *nspac* parameter. A **NULL** value in that location indicates that the caller doesn't wish to have the namespace returned. The *nspac* array must be at least of size one more than **PMIX_MAX_NSPLEN**.

By default, the spawned processes will be PMIx “connected” to the parent process upon successful launch (see Section 12.3 for details). This includes that (a) the parent process will be given a copy of the new job's information so it can query job-level info without incurring any communication penalties, (b) newly spawned child processes will receive a copy of the parent processes job-level info, and (c) both the parent process and members of the child job will receive notification of errors from processes in their combined assemblage.

Advice to users

Behavior of individual resource managers may differ, but it is expected that failure of any application process to start will result in termination/cleanup of all processes in the newly spawned job and return of an error code to the caller.

Advice to PMIx library implementers

Tools may utilize **PMIx_Spawn** to start intermediate launchers as described in Section 18.2.2. For times where the tool is not attached to a PMIx server, internal support for fork/exec of the specified applications would allow the tool to maintain a single code path for both the connected and disconnected cases. Inclusion of such support is recommended, but not required.

12.2.2 PMIx_Spawn_nb

Summary

Nonblocking version of the **PMIx_Spawn** routine.

Format

C

```
pmix_status_t
PMIx_Spawn_nb(const pmix_info_t job_info[], size_t ninfo,
              const pmix_app_t apps[], size_t napps,
              pmix_spawn_cbfunc_t cbfunc, void *cbdata)
```

C

IN **job_info**
Array of info structures (array of handles)

IN **ninfo**
Number of elements in the *job_info* array (integer)

IN **apps**
Array of [pmix_app_t](#) structures (array of handles)

IN **cbfunc**
Callback function [pmix_spawn_cbfunc_t](#) (function reference)

IN **cbdata**
Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when [PMIX_SUCCESS](#) is returned.

Returns [PMIX_SUCCESS](#) or one of the following error codes when the condition described occurs:

- [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing.

Host environments are required to support the following attributes when present in either the *job_info* or the *info* array of an element of the *apps* array:

PMIX_WDIR "pmix.wdir" (**char***)

Working directory for spawned processes.

PMIX_SET_SESSION_CWD "pmix.ssn cwd" (**bool**)

Set the current working directory to the session working directory assigned by the RM - can be assigned to the entire job (by including attribute in the *job_info* array) or on a per-application basis in the *info* array for each [pmix_app_t](#).

PMIX_PREFIX "pmix.prefix" (**char***)


1 Prefix to use for starting spawned processes - i.e., the directory where the executables can be
2 found.

3 **PMIX_HOST** "pmix.host" (char*)

4 Comma-delimited list of hosts to use for spawned processes.

5 **PMIX_HOSTFILE** "pmix.hostfile" (char*)

6 Hostfile to use for spawned processes.



7 The following attributes are optional for host environments that support this operation:

8 **PMIX_ADD_HOSTFILE** "pmix.addhostfile" (char*)

9 Hostfile containing hosts to add to existing allocation.

10 **PMIX_ADD_HOST** "pmix.addhost" (char*)

11 Comma-delimited list of hosts to add to the allocation.

12 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

13 Preload executables onto nodes prior to executing launch procedure.

14 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

15 Comma-delimited list of files to pre-position on nodes prior to executing launch procedure.

16 **PMIX_PERSONALITY** "pmix.pers" (char*)

17 Name of personality corresponding to programming model used by application - supported
18 values depend upon PMIx implementation.

19 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)

20 Display process mapping upon spawn.

21 **PMIX_PPR** "pmix.ppr" (char*)

22 Number of processes to spawn on each identified resource.

23 **PMIX_MAPBY** "pmix.mapby" (char*)

24 Process mapping policy - when accessed using **PMIx_Get**, use the
25 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
26 provided namespace. Supported values are launcher specific.

27 **PMIX_RANKBY** "pmix.rankby" (char*)

28 Process ranking policy - when accessed using **PMIx_Get**, use the
29 **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the
30 provided namespace. Supported values are launcher specific.

31 **PMIX_BINDTO** "pmix.bindto" (char*)

32 Process binding policy - when accessed using **PMIx_Get**, use the
33 **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the
34 provided namespace. Supported values are launcher specific.

1 **PMIX_STDIN_TGT** "pmix.stdin" (uint32_t)
 2 Spawned process rank that is to receive any forwarded **stdin**.

3 **PMIX_TAG_OUTPUT** "pmix.tagout" (bool)
 4 Tag **stdout/stderr** with the identity of the source process - can be assigned to the entire
 5 job (by including attribute in the *job_info* array) or on a per-application basis in the *info*
 6 array for each **pmix_app_t**.

7 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)
 8 Timestamp output - can be assigned to the entire job (by including attribute in the *job_info*
 9 array) or on a per-application basis in the *info* array for each **pmix_app_t**.

10 **PMIX_MERGE_STDERR_STDOUT** "pmix.mergeerrout" (bool)
 11 Merge **stdout** and **stderr** streams - can be assigned to the entire job (by including
 12 attribute in the *job_info* array) or on a per-application basis in the *info* array for each
 13 **pmix_app_t**.

14 **PMIX_OUTPUT_TO_FILE** "pmix.outfile" (char*)
 15 Direct output (both **stdout** and **stderr**) into files of form "<filename>.rank" - can be
 16 assigned to the entire job (by including attribute in the *job_info* array) or on a per-application
 17 basis in the *info* array for each **pmix_app_t**.

18 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)
 19 Mark the **argv** with the rank of the process.

20 **PMIX_CPUS_PER_PROC** "pmix.cputerproc" (uint32_t)
 21 Number of PUs to assign to each rank - when accessed using **PMIx_Get**, use the
 22 **PMIX_RANK_WILDCARD** value for the rank to discover the PUs/process assigned to the
 23 provided namespace.

24 **PMIX_NO_PROCS_ON_HEAD** "pmix.nolocal" (bool)
 25 Do not place processes on the head node.

26 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)
 27 Do not oversubscribe the nodes - i.e., do not place more processes than allocated slots on a
 28 node.

29 **PMIX_REPORT_BINDINGS** "pmix.repbinding" (bool)
 30 Report bindings of the individual processes.

31 **PMIX_CPU_LIST** "pmix.cpulists" (char*)
 32 List of PUs to use for this job - when accessed using **PMIx_Get**, use the
 33 **PMIX_RANK_WILDCARD** value for the rank to discover the PU list used for the provided
 34 namespace.

35 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)
 36 Application supports recoverable operations.

37 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)

1 Application is continuous, all failed processes should be immediately restarted.

2 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)

3 Maximum number of times to restart a process - when accessed using **PMIx_Get**, use the
4 **PMIX_RANK_WILDCARD** value for the rank to discover the max restarts for the provided
5 namespace.

6 **PMIX_SET_ENVAR** "pmix.envar.set" (pmix_envar_t*)

7 Set the envar to the given value, overwriting any pre-existing one

8 **PMIX_UNSET_ENVAR** "pmix.envar.unset" (char*)

9 Unset the environment variable specified in the string.

10 **PMIX_ADD_ENVAR** "pmix.envar.add" (pmix_envar_t*)

11 Add the environment variable, but do not overwrite any pre-existing one

12 **PMIX_PREPEND_ENVAR** "pmix.envar.prepnd" (pmix_envar_t*)

13 Prepend the given value to the specified environmental value using the given separator
14 character, creating the variable if it doesn't already exist

15 **PMIX_APPEND_ENVAR** "pmix.envar.appnd" (pmix_envar_t*)

16 Append the given value to the specified environmental value using the given separator
17 character, creating the variable if it doesn't already exist

18 **PMIX_FIRST_ENVAR** "pmix.envar.first" (pmix_envar_t*)

19 Ensure the given value appears first in the specified envar using the separator character,
20 creating the envar if it doesn't already exist

21 **PMIX_ALLOC_QUEUE** "pmix.alloc.queue" (char*)

22 Name of the WLM queue to which the allocation request is to be directed, or the queue being
23 referenced in a query.

24 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)

25 Total session time (in seconds) being requested in an allocation request.

26 **PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)

27 The number of nodes being requested in an allocation request.

28 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)

29 Regular expression of the specific nodes being requested in an allocation request.

30 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)

31 Number of PUs being requested in an allocation request.

32 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)

33 Regular expression of the number of PUs for each node being requested in an allocation
34 request.

35 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)

36 Regular expression of the specific PUs being requested in an allocation request.

1 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
2 Number of Megabytes[base2] of memory (per process) being requested in an allocation
3 request.

4 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)
5 Fabric bandwidth (in Megabits[base2]/sec) for the job being requested in an allocation
6 request.

7 **PMIX_ALLOC_FABRIC_QOS** "pmix.alloc.netqos" (char*)
8 Fabric quality of service level for the job being requested in an allocation request.

9 **PMIX_ALLOC_FABRIC_TYPE** "pmix.alloc.nettype" (char*)
10 Type of desired transport (e.g., "tcp", "udp") being requested in an allocation request.

11 **PMIX_ALLOC_FABRIC_PLANE** "pmix.alloc.netplane" (char*)
12 ID string for the *fabric plane* to be used for the requested allocation.

13 **PMIX_ALLOC_FABRIC_ENDPTS** "pmix.alloc.endpts" (size_t)
14 Number of endpoints to allocate per *process* in the job.

15 **PMIX_ALLOC_FABRIC_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)
16 Number of endpoints to allocate per *node* for the job.

17 **PMIX_COSPAWN_APP** "pmix.cospawn" (bool)
18 Designated application is to be spawned as a disconnected job - i.e., the launcher shall not
19 include the application in any of the job-level values (e.g., **PMIX_RANK** within the job)
20 provided to any other application process generated by the same spawn request. Typically
21 used to cospawn debugger daemons alongside an application.

22 **PMIX_SPAWN_TOOL** "pmix.spwn.tool" (bool)
23 Indicate that the job being spawned is a tool.

24 **PMIX_EVENT_SILENT_TERMINATION** "pmix.avsilentterm" (bool)
25 Do not generate an event when this job normally terminates.

26 **PMIX_ENVARS_HARVESTED** "pmix.evar.hvstd" (bool)
27 Environmental parameters have been harvested by the spawn requestor - the server does not
28 need to harvest them.

29 **PMIX_JOB_TIMEOUT** "pmix.job.time" (int)
30 Time in seconds before the spawned job should time out and be terminated (0 => infinite),
31 defined as the total runtime of the job (equivalent to the walltime limit of typical batch
32 schedulers).

33 **PMIX_SPAWN_TIMEOUT** "pmix.sp.time" (int)
34 Time in seconds before spawn operation should time out (0 => infinite). Logically
35 equivalent to passing the **PMIX_TIMEOUT** attribute to the **PMIx_Spawn** API, it is
36 provided as a separate attribute to distinguish it from the **PMIX_JOB_TIMEOUT** attribute

▲-----▲

Description

Nonblocking version of the `PMIx_Spawn` routine. The provided callback function will be executed upon successful start of *all* specified application processes.

Advice to users

Behavior of individual resource managers may differ, but it is expected that failure of any application process to start will result in termination/cleanup of all processes in the newly spawned job and return of an error code to the caller.

12.2.3 Spawn-specific constants

In addition to the generic error constants, the following spawn-specific error constants may be returned by the spawn APIs:

<code>PMIX_ERR_JOB_ALLOC_FAILED</code>	-188	The job request could not be executed due to failure to obtain the specified allocation
<code>PMIX_ERR_JOB_APP_NOT_EXECUTABLE</code>	-177	The specified application executable either could not be found, or lacks execution privileges.
<code>PMIX_ERR_JOB_NO_EXE_SPECIFIED</code>	-178	The job request did not specify an executable.
<code>PMIX_ERR_JOB_FAILED_TO_MAP</code>	-179	The launcher was unable to map the processes for the specified job request.
<code>PMIX_ERR_JOB_FAILED_TO_LAUNCH</code>	-181	One or more processes in the job request failed to launch
<i>Provisional</i> <code>PMIX_ERR_JOB_EXE_NOT_FOUND</code>		Specified executable not found
<i>Provisional</i> <code>PMIX_ERR_JOB_INSUFFICIENT_RESOURCES</code>		Insufficient resources to spawn job
<i>Provisional</i> <code>PMIX_ERR_JOB_SYS_OP_FAILED</code>		System library operation failed
<i>Provisional</i> <code>PMIX_ERR_JOB_WDIR_NOT_FOUND</code>		Specified working directory not found

1 12.2.4 Spawn attributes

2 Attributes used to describe [PMIx_Spawn](#) behavior - they are values passed to the [PMIx_Spawn](#)
3 API and therefore are not accessed using the [PMIx_Get](#) APIs when used in that context. However,
4 some of the attributes defined in this section can be provided by the host environment for other
5 purposes - e.g., the host might provide the [PMIX_MAPBY](#) attribute in the job-related information so
6 that an application can use [PMIx_Get](#) to discover the mapping used for determining process
7 locations. Multi-use attributes and their respective access reference rank are denoted below.

8 **PMIX_PERSONALITY** "pmix.pers" (char*)

9 Name of personality corresponding to programming model used by application - supported
10 values depend upon PMIx implementation.

11 **PMIX_HOST** "pmix.host" (char*)

12 Comma-delimited list of hosts to use for spawned processes.

13 **PMIX_HOSTFILE** "pmix.hostfile" (char*)

14 Hostfile to use for spawned processes.

15 **PMIX_ADD_HOST** "pmix.addhost" (char*)

16 Comma-delimited list of hosts to add to the allocation.

17 **PMIX_ADD_HOSTFILE** "pmix.addhostfile" (char*)

18 Hostfile containing hosts to add to existing allocation.

19 **PMIX_PREFIX** "pmix.prefix" (char*)

20 Prefix to use for starting spawned processes - i.e., the directory where the executables can be
21 found.

22 **PMIX_WDIR** "pmix.wdir" (char*)

23 Working directory for spawned processes.

24 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)

25 Display process mapping upon spawn.

26 **PMIX_PPR** "pmix.ppr" (char*)

27 Number of processes to spawn on each identified resource.

28 **PMIX_MAPBY** "pmix.mapby" (char*)

29 Process mapping policy - when accessed using [PMIx_Get](#), use the
30 [PMIX_RANK_WILDCARD](#) value for the rank to discover the mapping policy used for the
31 provided namespace. Supported values are launcher specific.

32 **PMIX_RANKBY** "pmix.rankby" (char*)

33 Process ranking policy - when accessed using [PMIx_Get](#), use the
34 [PMIX_RANK_WILDCARD](#) value for the rank to discover the ranking algorithm used for the
35 provided namespace. Supported values are launcher specific.

36 **PMIX_BINDTO** "pmix.bindto" (char*)

37 Process binding policy - when accessed using [PMIx_Get](#), use the
38 [PMIX_RANK_WILDCARD](#) value for the rank to discover the binding policy used for the
39 provided namespace. Supported values are launcher specific.

40 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

41 Preload executables onto nodes prior to executing launch procedure.

42 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

1 Comma-delimited list of files to pre-position on nodes prior to executing launch procedure.

2 **PMIX_STDIN_TGT** "pmix.stdin" (uint32_t)

3 Spawned process rank that is to receive any forwarded **stdin**.

4 **PMIX_SET_SESSION_CWD** "pmix.ssn cwd" (bool)

5 Set the current working directory to the session working directory assigned by the RM - can
6 be assigned to the entire job (by including attribute in the *job_info* array) or on a
7 per-application basis in the *info* array for each **pmix_app_t**.

8 **PMIX_TAG_OUTPUT** "pmix.tagout" (bool)

9 Tag **stdout/stderr** with the identity of the source process - can be assigned to the entire
10 job (by including attribute in the *job_info* array) or on a per-application basis in the *info*
11 array for each **pmix_app_t**.

12 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)

13 Timestamp output - can be assigned to the entire job (by including attribute in the *job_info*
14 array) or on a per-application basis in the *info* array for each **pmix_app_t**.

15 **PMIX_MERGE_STDERR_STDOUT** "pmix.mergeerrout" (bool)

16 Merge **stdout** and **stderr** streams - can be assigned to the entire job (by including
17 attribute in the *job_info* array) or on a per-application basis in the *info* array for each
18 **pmix_app_t**.

19 **PMIX_OUTPUT_TO_FILE** "pmix.outfile" (char*)

20 Direct output (both **stdout** and **stderr**) into files of form "<filename>.rank" - can be
21 assigned to the entire job (by including attribute in the *job_info* array) or on a per-application
22 basis in the *info* array for each **pmix_app_t**.

23 **PMIX_OUTPUT_TO_DIRECTORY** "pmix.outdir" (char*)

24 Direct output into files of form "<directory>/<jobid>/rank.<rank>/
25 **stdout[err]**" - can be assigned to the entire job (by including attribute in the *job_info*
26 array) or on a per-application basis in the *info* array for each **pmix_app_t**.

27 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)

28 Mark the **argv** with the rank of the process.

29 **PMIX_CPUS_PER_PROC** "pmix.cputerproc" (uint32_t)

30 Number of PUs to assign to each rank - when accessed using **PMIx_Get**, use the
31 **PMIX_RANK_WILDCARD** value for the rank to discover the PUs/process assigned to the
32 provided namespace.

33 **PMIX_NO_PROCS_ON_HEAD** "pmix.nolocal" (bool)

34 Do not place processes on the head node.

35 **PMIX_NO_OVERSUBSCRIBE** "pmix.noover" (bool)

36 Do not oversubscribe the nodes - i.e., do not place more processes than allocated slots on a
37 node.

38 **PMIX_REPORT_BINDINGS** "pmix.repbinding" (bool)

39 Report bindings of the individual processes.

40 **PMIX_CPU_LIST** "pmix.cpulist" (char*)

41 List of PUs to use for this job - when accessed using **PMIx_Get**, use the
42 **PMIX_RANK_WILDCARD** value for the rank to discover the PU list used for the provided
43 namespace.

1 **PMIX_JOB_RECOVERABLE** "pmix.recover" (bool)
2 Application supports recoverable operations.

3 **PMIX_JOB_CONTINUOUS** "pmix.continuous" (bool)
4 Application is continuous, all failed processes should be immediately restarted.

5 **PMIX_MAX_RESTARTS** "pmix.maxrestarts" (uint32_t)
6 Maximum number of times to restart a process - when accessed using [PMIx_Get](#), use the
7 [PMIX_RANK_WILDCARD](#) value for the rank to discover the max restarts for the provided
8 namespace.

9 **PMIX_SPAWN_TOOL** "pmix.spwn.tool" (bool)
10 Indicate that the job being spawned is a tool.

11 **PMIX_TIMEOUT_STACKTRACES** "pmix.tim.stack" (bool)
12 Include process stacktraces in timeout report from a job.

13 **PMIX_TIMEOUT_REPORT_STATE** "pmix.tim.state" (bool)
14 Report process states in timeout report from a job.

15 **PMIX_NOTIFY_JOB_EVENTS** "pmix.note.jev" (bool)
16 Requests that the launcher generate the [PMIX_EVENT_JOB_START](#),
17 [PMIX_LAUNCH_COMPLETE](#), and [PMIX_EVENT_JOB_END](#) events. Each event is to
18 include at least the namespace of the corresponding job and a [PMIX_EVENT_TIMESTAMP](#)
19 indicating the time the event occurred. Note that the requester must register for these
20 individual events, or capture and process them by registering a default event handler instead
21 of individual handlers and then process the events based on the returned status code.
22 Another common method is to register one event handler for all job-related events, with a
23 separate handler for non-job events - see [PMIx_Register_event_handler](#) for details.

24 **PMIX_NOTIFY_COMPLETION** "pmix.notecomp" (bool)
25 Requests that the launcher generate the [PMIX_EVENT_JOB_END](#) event for normal or
26 abnormal termination of the spawned job. The event shall include the returned status code
27 ([PMIX_JOB_TERM_STATUS](#)) for the corresponding job; the identity ([PMIX_PROCID](#))
28 and exit status ([PMIX_EXIT_CODE](#)) of the first failed process, if applicable; and a
29 [PMIX_EVENT_TIMESTAMP](#) indicating the time the termination occurred. Note that the
30 requester must register for the event or capture and process it within a default event handler.

31 **PMIX_NOTIFY_PROC_TERMINATION** "pmix.noteproc" (bool)
32 Requests that the launcher generate the [PMIX_EVENT_PROC_TERMINATED](#) event
33 whenever a process either normally or abnormally terminates.

34 **PMIX_NOTIFY_PROC_ABNORMAL_TERMINATION** "pmix.noteabproc" (bool)
35 Requests that the launcher generate the [PMIX_EVENT_PROC_TERMINATED](#) event only
36 when a process abnormally terminates.

37 **PMIX_LOG_PROC_TERMINATION** "pmix.logproc" (bool)
38 Requests that the launcher log the [PMIX_EVENT_PROC_TERMINATED](#) event whenever a
39 process either normally or abnormally terminates.

40 **PMIX_LOG_PROC_ABNORMAL_TERMINATION** "pmix.logabproc" (bool)
41 Requests that the launcher log the [PMIX_EVENT_PROC_TERMINATED](#) event only when a
42 process abnormally terminates.

43 **PMIX_LOG_JOB_EVENTS** "pmix.log.jev" (bool)

1 Requests that the launcher log the `PMIX_EVENT_JOB_START`,
2 `PMIX_LAUNCH_COMPLETE`, and `PMIX_EVENT_JOB_END` events using `PMIx_Log`,
3 subject to the logging attributes of Section 13.4.3.

4 **PMIX_LOG_COMPLETION** "`pmix.logcomp`" (bool)

5 Requests that the launcher log the `PMIX_EVENT_JOB_END` event for normal or abnormal
6 termination of the spawned job using `PMIx_Log`, subject to the logging attributes of
7 Section 13.4.3. The event shall include the returned status code
8 (`PMIX_JOB_TERM_STATUS`) for the corresponding job; the identity (`PMIX_PROCID`)
9 and exit status (`PMIX_EXIT_CODE`) of the first failed process, if applicable; and a
10 `PMIX_EVENT_TIMESTAMP` indicating the time the termination occurred.

11 **PMIX_EVENT_SILENT_TERMINATION** "`pmix.evsilentterm`" (bool)

12 Do not generate an event when this job normally terminates.

13 *Provisional* **PMIX_ENVARS_HARVESTED** "`pmix.avar.hvstd`" (bool)

14 Environmental parameters have been harvested by the spawn requestor - the server does not
15 need to harvest them.

16 *Provisional* **PMIX_JOB_TIMEOUT** "`pmix.job.time`" (int)

17 Time in seconds before the spawned job should time out and be terminated (0 => infinite),
18 defined as the total runtime of the job (equivalent to the walltime limit of typical batch
19 schedulers).

20 *Provisional* **PMIX_SPAWN_TIMEOUT** "`pmix.sp.time`" (int)

21 Time in seconds before spawn operation should time out (0 => infinite). Logically
22 equivalent to passing the `PMIX_TIMEOUT` attribute to the `PMIx_Spawn` API, it is
23 provided as a separate attribute to distinguish it from the `PMIX_JOB_TIMEOUT` attribute

24 Attributes used to adjust remote environment variables prior to spawning the specified application
25 processes.

26 **PMIX_SET_ENVAR** "`pmix.avar.set`" (pmix_avar_t*)

27 Set the envar to the given value, overwriting any pre-existing one

28 **PMIX_UNSET_ENVAR** "`pmix.avar.unset`" (char*)

29 Unset the environment variable specified in the string.

30 **PMIX_ADD_ENVAR** "`pmix.avar.add`" (pmix_avar_t*)

31 Add the environment variable, but do not overwrite any pre-existing one

32 **PMIX_PREPEND_ENVAR** "`pmix.avar.prepnd`" (pmix_avar_t*)

33 Prepend the given value to the specified environmental value using the given separator
34 character, creating the variable if it doesn't already exist

35 **PMIX_APPEND_ENVAR** "`pmix.avar.appnd`" (pmix_avar_t*)

36 Append the given value to the specified environmental value using the given separator
37 character, creating the variable if it doesn't already exist

38 **PMIX_FIRST_ENVAR** "`pmix.avar.first`" (pmix_avar_t*)

39 Ensure the given value appears first in the specified envar using the separator character,
40 creating the envar if it doesn't already exist

1 12.2.5 Application Structure

2 The `pmix_app_t` structure describes the application context for the `PMIx_Spawn` and
3 `PMIx_Spawn_nb` operations.

PMIx v1.0

C

```
4 typedef struct pmix_app {  
5     /** Executable */  
6     char *cmd;  
7     /** Argument set, NULL terminated */  
8     char **argv;  
9     /** Environment set, NULL terminated */  
10    char **env;  
11    /** Current working directory */  
12    char *cwd;  
13    /** Maximum processes with this profile */  
14    int maxprocs;  
15    /** Array of info keys describing this application*/  
16    pmix_info_t *info;  
17    /** Number of info keys in 'info' array */  
18    size_t ninfo;  
19 } pmix_app_t;
```

C

20 12.2.5.1 App structure support macros

21 The following macros are provided to support the `pmix_app_t` structure.

22 **Static initializer for the app structure**

23 *Provisional* Provide a static initializer for the `pmix_app_t` fields.

PMIx v4.2

C

24 `PMIX_APP_STATIC_INIT`

C

25 **Initialize the app structure**

26 Initialize the `pmix_app_t` fields

PMIx v1.0

C

27 `PMIX_APP_CONSTRUCT (m)`

C

28 **IN** m

29 Pointer to the structure to be initialized (pointer to `pmix_app_t`)

1 **Destruct the app structure**

2 Destruct the `pmix_app_t` fields



3 **PMIX_APP_DESTRUCT (m)**



4 **IN** m

5 Pointer to the structure to be destructed (pointer to `pmix_app_t`)

6 **Create an app array**

7 Allocate and initialize an array of `pmix_app_t` structures

PMIx v1.0



8 **PMIX_APP_CREATE (m, n)**



9 **INOUT** m

10 Address where the pointer to the array of `pmix_app_t` structures shall be stored (handle)

11 **IN** n

12 Number of structures to be allocated (`size_t`)

13 **Free an app structure**

14 Release a `pmix_app_t` structure

PMIx v4.0



15 **PMIX_APP_RELEASE (m)**



16 **IN** m

17 Pointer to a `pmix_app_t` structure (handle)

18 **Free an app array**

19 Release an array of `pmix_app_t` structures

PMIx v1.0



20 **PMIX_APP_FREE (m, n)**



21 **IN** m

22 Pointer to the array of `pmix_app_t` structures (handle)

23 **IN** n

24 Number of structures in the array (`size_t`)

1 Create the info array of application directives

2 Create an array of `pmix_info_t` structures for passing application-level directives, updating the
3 `ninfo` field of the `pmix_app_t` structure.

```
▼ _____ C _____ ▼
```

4 `PMIX_APP_INFO_CREATE(m, n)`

```
▲ _____ C _____ ▲
```

5 **IN** `m`

6 Pointer to the `pmix_app_t` structure (handle)

7 **IN** `n`

8 Number of directives to be allocated (`size_t`)

9 12.2.5.2 Spawn Callback Function

10 Summary

11 The `pmix_spawn_cbfunc_t` is used on the PMIx client side by `PMIx_Spawn_nb` and on the
12 PMIx server side by `pmix_server_spawn_fn_t`.

PMIx v1.0 ▼ _____ C _____ ▼

```
13 typedef void (*pmix_spawn_cbfunc_t)  
14     (pmix_status_t status,  
15      pmix_nspace_t nspace, void *cbdata);
```

```
▲ _____ C _____ ▲
```

16 **IN** `status`

17 Status associated with the operation (handle)

18 **IN** `nspace`

19 Namespace string (`pmix_nspace_t`)

20 **IN** `cbdata`

21 Callback data passed to original API call (memory reference)

22 Description

23 The callback will be executed upon launch of the specified applications in `PMIx_Spawn_nb`, or
24 upon failure to launch any of them.

25 The `status` of the callback will indicate whether or not the spawn succeeded. The `nspace` of the
26 spawned processes will be returned, along with any provided callback data. Note that the returned
27 `nspace` value will not be protected upon return from the callback function, so the receiver must
28 copy it if it needs to be retained.

1 12.3 Connecting and Disconnecting Processes

2 This section defines functions to connect and disconnect processes in two or more separate PMIx
3 namespaces. The PMIx definition of *connected* solely implies that the host environment should
4 treat the failure of any process in the assemblage as a reportable event, taking action on the
5 assemblage as if it were a single application. For example, if the environment defaults (in the
6 absence of any application directives) to terminating an application upon failure of any process in
7 that application, then the environment should terminate all processes in the connected assemblage
8 upon failure of any member.

9 The host environment may choose to assign a new namespace to the connected assemblage and/or
10 assign new ranks for its members for its own internal tracking purposes. However, it is not required
11 to communicate such assignments to the participants (e.g., in response to an appropriate call to
12 [PMIx_Query_info_nb](#)). The host environment is required to generate a
13 [PMIX_ERR_PROC_TERM_WO_SYNC](#) event should any process in the assemblage terminate or
14 call [PMIx_Finalize](#) without first *disconnecting* from the assemblage. If the job including the
15 process is terminated as a result of that action, then the host environment is required to also
16 generate the [PMIX_ERR_JOB_TERM_WO_SYNC](#) for all jobs that were terminated as a result.

Advice to PMIx server hosts

17 The *connect* operation does not require the exchange of job-level information nor the inclusion of
18 information posted by participating processes via [PMIx_Put](#). Indeed, the callback function
19 utilized in [pmix_server_connect_fn_t](#) cannot pass information back into the PMIx server
20 library. However, host environments are advised that collecting such information at the
21 participating daemons represents an optimization opportunity as participating processes are likely
22 to request such information after the connect operation completes.

Advice to users

23 Attempting to *connect* processes solely within the same namespace is essentially a *no-op* operation.
24 While not explicitly prohibited, users are advised that a PMIx implementation or host environment
25 may return an error in such cases.

26 Neither the PMIx implementation nor host environment are required to provide any tracking
27 support for the assemblage. Thus, the application is responsible for maintaining the membership
28 list of the assemblage.

29 12.3.1 PMIx_Connect

30 Summary

31 Connect namespaces.

Format

C

```
pmix_status_t
PMIx_Connect(const pmix_proc_t procs[], size_t nprocs,
             const pmix_info_t info[], size_t ninfo)
```

C

IN `procs`
Array of proc structures (array of handles)

IN `nprocs`
Number of elements in the *procs* array (integer)

IN `info`
Array of info structures (array of handles)

IN `ninfo`
Number of elements in the *info* array (integer)

Returns **PMIX_SUCCESS** or a negative value indicating the error.

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

The following attributes are optional for PMIx implementations:

PMIX_ALL_CLONES_PARTICIPATE "`pmix.clone.part`" (**bool**)
All *clones* of the calling process must participate in the collective operation.

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "`pmix.timeout`" (**int**)
Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Record the processes specified by the *procs* array as *connected* as per the PMIx definition. The function will return once all processes identified in *procs* have called either **PMIx_Connect** or its non-blocking version, *and* the host environment has completed any supporting operations required to meet the terms of the PMIx definition of *connected* processes.

A process can only engage in one connect operation involving the identical *procs* array at a time. However, a process can be simultaneously engaged in multiple connect operations, each involving a different *procs* array.

As in the case of the **PMIx_Fence** operation, the *info* array can be used to pass user-level directives regarding timeout constraints and other options available from the host RM.

Advice to users

All processes engaged in a given **PMIx_Connect** operation must provide the identical *procs* array as ordering of entries in the array and the method by which those processes are identified (e.g., use of **PMIX_RANK_WILDCARD** versus listing the individual processes) *may* impact the host environment's algorithm for uniquely identifying an operation.

Advice to PMIx library implementers

PMIx_Connect and its non-blocking form are both *collective* operations. Accordingly, the PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective.

12.3.2 PMIx_Connect_nb

Summary

Nonblocking **PMIx_Connect_nb** routine.

Format

C

```
pmix_status_t
PMIx_Connect_nb(const pmix_proc_t procs[], size_t nprocs,
               const pmix_info_t info[], size_t ninfo,
               pmix_op_cbfunc_t cbfunc, void *cbdata)
```

C

- IN** **procs**
Array of proc structures (array of handles)
- IN** **nprocs**
Number of elements in the *procs* array (integer)
- IN** **info**
Array of info structures (array of handles)
- IN** **ninfo**
Number of elements in the *info* array (integer)
- IN** **cbfunc**
Callback function [pmix_op_cbfunc_t](#) (function reference)
- IN** **cbdata**
Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

If none of the above return codes are appropriate, then an implementation must return either a general **PMIX** error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

PMIX libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

The following attributes are optional for PMIx implementations:

PMIX_ALL_CLONES_PARTICIPATE "pmix.clone.part" (bool)

All *clones* of the calling process must participate in the collective operation.

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Nonblocking version of **PMIx_Connect**. The callback function is called once all processes identified in *procs* have called either **PMIx_Connect** or its non-blocking version, *and* the host environment has completed any supporting operations required to meet the terms of the PMIx definition of *connected* processes. See the advice provided in the description for **PMIx_Connect** for more information.

12.3.3 PMIx_Disconnect

Summary

Disconnect a previously connected set of processes.

Format

PMIx v1.0

```
pmix_status_t
PMIx_Disconnect(const pmix_proc_t procs[], size_t nprocs,
                const pmix_info_t info[], size_t ninfo);
```

IN **procs**
Array of proc structures (array of handles)

IN **nprocs**
Number of elements in the *procs* array (integer)

IN **info**
Array of info structures (array of handles)

IN **ninfo**
Number of elements in the *info* array (integer)

Returns PMIX_SUCCESS or one of the following error codes when the condition described occurs:

- 1 • the **PMIX_ERR_INVALID_OPERATION** error indicating that the specified set of *procs* was not
2 previously *connected* via a call to **PMIx_Connect** or its non-blocking form.

3 If none of the above return codes are appropriate, then an implementation must return either a
4 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

5 PMIx libraries are not required to directly support any attributes for this function. However, any
6 provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

7 The following attributes are optional for PMIx implementations:

8 **PMIX_ALL_CLONES_PARTICIPATE** "**pmix.clone.part**" (**bool**)

9 All *clones* of the calling process must participate in the collective operation.

10 The following attributes are optional for host environments that support this operation:

11 **PMIX_TIMEOUT** "**pmix.timeout**" (**int**)

12 Time in seconds before the specified operation should time out (zero indicating infinite) and
13 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
14 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

15 Disconnect a previously connected set of processes. The function will return once all processes
16 identified in *procs* have called either **PMIx_Disconnect** or its non-blocking version, *and* the
17 host environment has completed any required supporting operations.
18

19 A process can only engage in one disconnect operation involving the identical *procs* array at a time.
20 However, a process can be simultaneously engaged in multiple disconnect operations, each
21 involving a different *procs* array.

22 As in the case of the **PMIx_Fence** operation, the *info* array can be used to pass user-level
23 directives regarding the algorithm to be used for any collective operation involved in the operation,
24 timeout constraints, and other options available from the host RM.

Advice to users

25 All processes engaged in a given **PMIx_Disconnect** operation must provide the identical *procs*
26 array as ordering of entries in the array and the method by which those processes are identified
27 (e.g., use of **PMIX_RANK_WILDCARD** versus listing the individual processes) *may* impact the host
28 environment's algorithm for uniquely identifying an operation.

Advice to PMIx library implementers

1 **PMIx_Disconnect** and its non-blocking form are both *collective* operations. Accordingly, the
2 PMIx server library is required to aggregate participation by local clients, passing the request to the
3 host environment once all local participants have executed the API.

Advice to PMIx server hosts

4 The host will receive a single call for each collective operation. The host will receive a single call
5 for each collective operation. It is the responsibility of the host to identify the nodes containing
6 participating processes, execute the collective across all participating nodes, and notify the local
7 PMIx server library upon completion of the global collective.

12.3.4 PMIx_Disconnect_nb

Summary

Nonblocking **PMIx_Disconnect** routine.

Format

PMIx v1.0

```
pmix_status_t  
PMIx_Disconnect_nb(const pmix_proc_t procs[], size_t nprocs,  
                  const pmix_info_t info[], size_t ninfo,  
                  pmix_op_cbfunc_t cbfunc, void *cbdata);
```

IN procs
Array of proc structures (array of handles)

IN nprocs
Number of elements in the *procs* array (integer)

IN info
Array of info structures (array of handles)

IN ninfo
Number of elements in the *info* array (integer)

IN cbfunc
Callback function **pmix_op_cbfunc_t** (function reference)

IN cbdata
Data to be passed to the callback function (memory reference)

1 A successful return indicates that the request is being processed and the result will be returned in
2 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
3 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

4 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 5 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
6 returned *success* - the *cbfunc* will *not* be called

7 If none of the above return codes are appropriate, then an implementation must return either a
8 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

9 PMIx libraries are not required to directly support any attributes for this function. However, any
10 provided attributes must be passed to the host SMS daemon for processing.

Optional Attributes

11 The following attributes are optional for PMIx implementations:

12 **PMIX_ALL_CLONES_PARTICIPATE** "pmix.clone.part" (bool)

13 All *clones* of the calling process must participate in the collective operation.

14 The following attributes are optional for host environments that support this operation:

15 **PMIX_TIMEOUT** "pmix.timeout" (int)

16 Time in seconds before the specified operation should time out (zero indicating infinite) and
17 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
18 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

19 Nonblocking **PMIx_Disconnect** routine. The callback function is called either:

- 21 • to return the **PMIX_ERR_INVALID_OPERATION** error indicating that the specified set of
22 *procs* was not previously *connected* via a call to **PMIx_Connect** or its non-blocking form;
- 23 • to return a PMIx error constant indicating that the operation failed; or
- 24 • once all processes identified in *procs* have called either **PMIx_Disconnect_nb** or its
25 blocking version, *and* the host environment has completed any required supporting operations.

26 See the advice provided in the description for **PMIx_Disconnect** for more information.

1 12.4 Process Locality

2 The relative locality of processes is often used to optimize their interactions with the hardware and
3 other processes. PMIx provides a means by which the host environment can communicate the
4 locality of a given process using the `PMIx_server_generate_locality_string` to
5 generate an abstracted representation of that value. This provides a human-readable format and
6 allows the client to parse the locality string with a method of its choice that may differ from the one
7 used by the server that generated it.

8 There are times, however, when relative locality and other PMIx-provided information doesn't
9 include some element required by the application. In these instances, the application may need
10 access to the full description of the local hardware topology. PMIx does not itself generate such
11 descriptions - there are multiple third-party libraries that fulfill that role. Instead, PMIx offers an
12 abstraction method by which users can obtain a pointer to the description. This transparently
13 enables support for different methods of sharing the topology between the host environment (which
14 may well have already generated it prior to local start of application processes) and the clients - e.g.,
15 through passing of a shared memory region.

16 12.4.1 PMIx_Load_topology

17 Summary

18 Load the local hardware topology description

19 Format

PMIx v4.0

```
20 pmix_status_t  
21 PMIx_Load_topology(pmix_topology_t *topo);
```

22 INOUT topo

23 Address of a `pmix_topology_t` structure where the topology information is to be loaded
24 (handle)

25 A successful return indicates that the *topo* was successfully loaded.

26 Returns `PMIX_SUCCESS` or a negative value indicating the error.

27 Description

28 Obtain a pointer to the topology description of the local node. If the *source* field of the provided
29 `pmix_topology_t` is set, then the PMIx library must return a description from the specified
30 implementation or else indicate that the implementation is not available by returning the
31 `PMIX_ERR_NOT_SUPPORTED` error constant.

32 The returned pointer may point to a shared memory region or an actual instance of the topology
33 description. In either case, the description shall be treated as a "read-only" object - attempts to
34 modify the object are likely to fail and return an error. The PMIx library is responsible for
35 performing any required cleanup when the client library finalizes.

Advice to users

It is the responsibility of the user to ensure that the *topo* argument is properly initialized prior to calling this API, and to check the returned *source* to verify that the returned topology description is compatible with the user's code.

12.4.2 PMIx_Get_relative_locality

Summary

Get the relative locality of two local processes given their locality strings.

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Get_relative_locality(const char *locality1,
                          const char *locality2,
                          pmix_locality_t *locality);
```

C

IN locality1

String returned by the [PMIx_server_generate_locality_string](#) API (handle)

IN locality2

String returned by the [PMIx_server_generate_locality_string](#) API (handle)

INOUT locality

Location where the relative locality bitmask is to be constructed (memory reference)

A successful return indicates that the *locality* was successfully loaded.

Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Description

Parse the locality strings of two processes (as returned by [PMIx_Get](#) using the [PMIX_LOCALITY_STRING](#) key) and set the appropriate [pmix_locality_t](#) locality bits in the provided memory location.

12.4.2.1 Topology description

The [pmix_topology_t](#) structure contains a (case-insensitive) string identifying the source of the topology (e.g., "hwloc") and a pointer to the corresponding implementation-specific topology description.

PMIx v4.0

C

```
typedef struct pmix_topology {
    char *source;
    void *topology;
} pmix_topoology_t;
```

C

1 12.4.2.2 Topology support macros

2 The following macros support the `pmix_topology_t` structure.

3 **Static initializer for the topology structure**

4 **Provisional** Provide a static initializer for the `pmix_topology_t` fields.

PMIx v4.2

```
▼ _____ C _____ ▼  
PMIX_TOPOLOGY_STATIC_INIT  
▲ _____ C _____ ▲
```

6 **Initialize the topology structure**

7 Initialize the `pmix_topology_t` fields to `NULL`

PMIx v4.0

```
▼ _____ C _____ ▼  
PMIX_TOPOLOGY_CONSTRUCT (m)  
▲ _____ C _____ ▲
```

9 **IN** m

10 Pointer to the structure to be initialized (pointer to `pmix_topology_t`)

11 **Destruct a topology structure**

12 **Summary**

13 Destruct a `pmix_topology_t` fields

14 **Format**

PMIx v4.2

```
▼ _____ C _____ ▼  
void  
PMIx_Topology_destruct (pmix_topology_t *topo);  
▲ _____ C _____ ▲
```

17 **IN** topo

18 Pointer to the structure to be destructed (pointer to `pmix_topology_t`)

19 **Description**

20 Release any memory storage held by the `pmix_topology_t` structure

21 **Create a topology array**

22 Allocate and initialize a `pmix_topology_t` array.

PMIx v4.0

```
▼ _____ C _____ ▼  
PMIX_TOPOLOGY_CREATE (m, n)  
▲ _____ C _____ ▲
```

24 **INOUT** m

25 Address where the pointer to the array of `pmix_topology_t` structures shall be stored
26 (handle)

27 **IN** n

28 Number of structures to be allocated (size_t)

1 12.4.2.3 Relative locality of two processes

2 The `pmix_locality_t` datatype is a `uint16_t` bitmask that defines the relative locality of
3 two processes on a node. The following constants represent specific bits in the mask and can be
4 used to test a locality value using standard bit-test methods.

5 **PMIX_LOCALITY_UNKNOWN** `0x0000` All bits are set to zero, indicating that the relative
6 locality of the two processes is unknown

7 **PMIX_LOCALITY_NONLOCAL** `0x0000` The two processes do not share any common
8 locations

9 **PMIX_LOCALITY_SHARE_HWTHREAD** `0x0001` The two processes share at least one
10 hardware thread

11 **PMIX_LOCALITY_SHARE_CORE** `0x0002` The two processes share at least one core

12 **PMIX_LOCALITY_SHARE_L1CACHE** `0x0004` The two processes share at least an L1
13 cache

14 **PMIX_LOCALITY_SHARE_L2CACHE** `0x0008` The two processes share at least an L2
15 cache

16 **PMIX_LOCALITY_SHARE_L3CACHE** `0x0010` The two processes share at least an L3
17 cache

18 **PMIX_LOCALITY_SHARE_PACKAGE** `0x0020` The two processes share at least a package

19 **PMIX_LOCALITY_SHARE_NUMA** `0x0040` The two processes share at least one
20 Non-Uniform Memory Access (NUMA) region

21 **PMIX_LOCALITY_SHARE_NODE** `0x4000` The two processes are executing on the same
22 node

23 Implementers and vendors may choose to extend these definitions as needed to describe a particular
24 system.

25 12.4.2.4 Locality keys

26 **PMIX_LOCALITY_STRING** `"pmix.locstr"` (`char*`)

27 String describing a process's bound location - referenced using the process's rank. The string
28 is prefixed by the implementation that created it (e.g., "hwloc") followed by a colon. The
29 remainder of the string represents the corresponding locality as expressed by the underlying
30 implementation. The entire string must be passed to `PMIx_Get_relative_locality`
31 for processing. Note that hosts are only required to provide locality strings for local client
32 processes - thus, a call to `PMIx_Get` for the locality string of a process that returns
33 `PMIX_ERR_NOT_FOUND` indicates that the process is not executing on the same node.

34 12.4.3 `PMIx_Parse_cpuset_string`

35 Summary

36 Parse the PU binding bitmap from its string representation.

Format

```
pmix_status_t
PMIx_Parse_cpuset_string(const char *cpuset_string,
                        pmix_cpuset_t *cpuset);
```

IN cpuset_string

String returned by the [PMIx_server_generate_cpuset_string](#) API (handle)

INOUT cpuset

Address of an object where the bitmap is to be stored (memory reference)

A successful return indicates that the *cpuset* was successfully loaded.

Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Description

Parse the string representation of the binding bitmap (as returned by [PMIx_Get](#) using the [PMIX_CPUSSET](#) key) and set the appropriate PU binding location information in the provided memory location.

12.4.4 PMIx_Get_cpuset

Summary

Get the PU binding bitmap of the current process.

Format

```
pmix_status_t
PMIx_Get_cpuset(pmix_cpuset_t *cpuset, pmix_bind_envelope_t ref);
```

INOUT cpuset

Address of an object where the bitmap is to be stored (memory reference)

IN ref

The binding envelope to be considered when formulating the bitmap
([pmix_bind_envelope_t](#))

A successful return indicates that the *cpuset* was successfully loaded.

Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Description

Obtain and set the appropriate PU binding location information in the provided memory location based on the specified binding envelope.

1 12.4.4.1 Binding envelope

2 The `pmix_bind_envelope_t` data type defines the envelope of threads within a possibly
3 multi-threaded process that are to be considered when getting the cpuset associated with the
4 process. Valid values include:

- 5 **PMIX_CPUBIND_PROCESS** 0 Use the location of all threads in the possibly multi-threaded
6 process.
7 **PMIX_CPUBIND_THREAD** 1 Use only the location of the thread calling the API.

8 12.4.5 PMIx_Compute_distances

9 Summary

10 Compute distances from specified process location to local devices.

11 Format

PMIx v4.0

C

12 `pmix_status_t`

```
13 PMIx_Compute_distances(pmix_topology_t *topo,  
14                       pmix_cpuset_t *cpuset,  
15                       pmix_info_t info[], size_t ninfo[],  
16                       pmix_device_distance_t *distances[],  
17                       size_t *ndist);
```

C

18 **IN topo**

19 Pointer to the topology description of the node where the process is located (**NULL** indicates
20 the local node) (`pmix_topology_t`)

21 **IN cpuset**

22 Pointer to the location of the process (`pmix_cpuset_t`)

23 **IN info**

24 Array of `pmix_info_t` describing the devices whose distance is to be computed (handle)

25 **IN ninfo**

26 Number of elements in *info* (integer)

27 **INOUT distances**

28 Pointer to an address where the array of `pmix_device_distance_t` structures
29 containing the distances from the caller to the specified devices is to be returned (handle)

30 **INOUT ndist**

31 Pointer to an address where the number of elements in the *distances* array is to be returned
32 (handle)

33 Returns **PMIX_SUCCESS** or a negative value indicating the error.

Description

Both the minimum and maximum distance fields in the elements of the array shall be filled with the respective distances between the current process location and the types of devices or specific device identified in the *info* directives. In the absence of directives, distances to all supported device types shall be returned.

Advice to users

A process whose threads are not all bound to the same location may return inconsistent results from calls to this API by different threads if the `PMIX_CPUBIND_THREAD` binding envelope was used when generating the *cpuset*.

12.4.6 PMIx_Compute_distances_nb

Summary

Compute distances from specified process location to local devices.

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Compute_distances_nb(pmix_topology_t *topo,
                          pmix_cpuset_t *cpuset,
                          pmix_info_t info[], size_t ninfo[],
                          pmix_device_dist_cbfunc_t cbfunc,
                          void *cbdata);
```

C

IN topo

Pointer to the topology description of the node where the process is located (**NULL** indicates the local node) (`pmix_topology_t`)

IN cpuset

Pointer to the location of the process (`pmix_cpuset_t`)

IN info

Array of `pmix_info_t` describing the devices whose distance is to be computed (handle)

IN ninfo

Number of elements in *info* (integer)

IN cbfunc

Callback function `pmix_info_cbfunc_t` (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

1 A successful return indicates that the request is being processed and the result will be returned in
2 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
3 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

4 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 5 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed
6 successfully - the *cbfunc* will *not* be called.

7 If none of the above return codes are appropriate, then an implementation must return either a
8 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

9 **Description**

10 Non-blocking form of the **PMIx_Compute_distances** API.

11 **12.4.7 Device Distance Callback Function**

12 **Summary**

13 The **pmix_device_dist_cbfunc_t** is used to return an array of device distances.

PMIx v4.0

```
14 typedef void (*pmix_device_dist_cbfunc_t)  
15     (pmix_status_t status,  
16      pmix_device_distance_t *dist,  
17      size_t ndist,  
18      void *cbdata,  
19      pmix_release_cbfunc_t release_fn,  
20      void *release_cbdata);
```

21 **IN status**

22 Status associated with the operation (**pmix_status_t**)

23 **IN dist**

24 Array of **pmix_device_distance_t** returned by the operation (pointer)

25 **IN ndist**

26 Number of elements in the *dist* array (**size_t**)

27 **IN cbdata**

28 Callback data passed to original API call (memory reference)

29 **IN release_fn**

30 Function to be called when done with the *dist* data (function pointer)

31 **IN release_cbdata**

32 Callback data to be passed to *release_fn* (memory reference)

33 **Description**

34 The *status* indicates if requested data was found or not. The array of
35 **pmix_device_distance_t** will contain the distance information.

1 12.4.8 Device type

2 The `pmix_device_type_t` is a `uint64_t` bitmask for identifying the type(s) whose
3 distances are being requested, or the type of a specific device being referenced (e.g., in a
4 `pmix_device_distance_t` object).

PMIx v1.0

```
▼ typedef uint16_t pmix_device_type_t; C ▲
```

6 The following constants can be used to set a variable of the type `pmix_device_type_t`.

7 **PMIX_DEVTYPE_UNKNOWN 0x00** The device is of an unknown type - will not be included
8 in returned device distances.
9 **PMIX_DEVTYPE_BLOCK 0x01** Operating system block device, or non-volatile memory
10 device (e.g., "sda" or "dax2.0" on Linux).
11 **PMIX_DEVTYPE_GPU 0x02** Operating system Graphics Processing Unit (GPU) device
12 (e.g., "card0" for a Linux Direct Rendering Manager (DRM) device).
13 **PMIX_DEVTYPE_NETWORK 0x04** Operating system network device (e.g., the "eth0"
14 interface on Linux).
15 **PMIX_DEVTYPE_OPENFABRICS 0x08** Operating system OpenFabrics device (e.g., an
16 "mlx4_0" InfiniBand Host Channel Adapter (HCA), or "hfi1_0" Omni-Path interface on
17 Linux).
18 **PMIX_DEVTYPE_DMA 0x10** Operating system Direct Memory Access (DMA) engine
19 device (e.g., the "dma0chan0" DMA channel on Linux).
20 **PMIX_DEVTYPE_COPROC 0x20** Operating system co-processor device (e.g., "mic0" for a
21 Xeon Phi on Linux, "opencl0d0" for a OpenCL device, or "cuda0" for a Compute Unified
22 Device Architecture (CUDA) device).

23 12.4.9 Device Distance Structure

24 The `pmix_device_distance_t` structure contains the minimum and maximum relative
25 distance from the caller to a given device.

PMIx v4.0

```
▼ typedef struct pmix_device_distance { C ▲  
26     char *uuid;  
27     char *osname;  
28     pmix_device_type_t type;  
29     uint16_t mindist;  
30     uint16_t maxdist;  
31     } pmix_device_distance_t;  
32
```


The *uuid* is a string identifier guaranteed to be unique within the cluster and is typically assembled from discovered device attributes (e.g., the Internet Protocol (IP) address of the device). The *osname* is the local operating system name of the device and is only unique to that node.

The two distance fields provide the minimum and maximum relative distance to the device from the specified location of the process, expressed as a 16-bit integer value where a smaller number indicates that this device is closer to the process than a device with a larger distance value. Note that relative distance values are not necessarily correlated to a physical property - e.g., a device at twice the distance from another device does not necessarily have twice the latency for communication with it.

Relative distances only apply to similar devices and cannot be used to compare devices of different types. Both minimum and maximum distances are provided to support cases where the process may be bound to more than one location, and the locations are at different distances from the device.

A relative distance value of **UINT16_MAX** indicates that the distance from the process to the device could not be provided. This may be due to lack of available information (e.g., the PMIx library not having access to device locations) or other factors.

12.4.10 Device distance support macros

The following macros are provided to support the `pmix_device_distance_t` structure.

Static initializer for the device distance structure

Provide a static initializer for the `pmix_device_distance_t` fields.

Provisional
PMIx v4.2

PMIX_DEVICE_DIST_STATIC_INIT

Initialize the device distance structure

Initialize the `pmix_device_distance_t` fields.

PMIx v4.0

PMIX_DEVICE_DIST_CONSTRUCT (m)

IN m

Pointer to the structure to be initialized (pointer to `pmix_device_distance_t`)

1 **Destruct the device distance structure**

2 Destruct the `pmix_device_distance_t` fields.



3 **PMIX_DEVICE_DIST_DESTRUCT (m)**



4 **IN** m

5 Pointer to the structure to be destructed (pointer to `pmix_device_distance_t`)

6 **Create an device distance array**

7 Allocate and initialize a `pmix_device_distance_t` array.

PMIx v4.0



8 **PMIX_DEVICE_DIST_CREATE (m, n)**



9 **INOUT** m

10 Address where the pointer to the array of `pmix_device_distance_t` structures shall be
11 stored (handle)

12 **IN** n

13 Number of structures to be allocated (`size_t`)

14 **Release an device distance array**

15 Release an array of `pmix_device_distance_t` structures.

PMIx v4.0



16 **PMIX_DEVICE_DIST_FREE (m, n)**



17 **IN** m

18 Pointer to the array of `pmix_device_distance_t` structures (handle)

19 **IN** n

20 Number of structures in the array (`size_t`)

1 12.4.11 Device distance attributes

2 The following attributes can be used to retrieve device distances from the PMIx data store. Note
3 that distances stored by the host environment are based on the process location at the time of start
4 of execution and may not reflect changes to location imposed by the process itself.

5 **PMIX_DEVICE_DISTANCES** "pmix.dev.dist" (**pmix_data_array_t**)

6 Return an array of **pmix_device_distance_t** containing the minimum and maximum
7 distances of the given process location to all devices of the specified type on the local node.

8 **PMIX_DEVICE_TYPE** "pmix.dev.type" (**pmix_device_type_t**)

9 Bitmask specifying the type(s) of device(s) whose information is being requested. Only used
10 as a directive/qualifier.

11 **PMIX_DEVICE_ID** "pmix.dev.id" (**string**)

12 System-wide Universally Unique Identifier (UUID) or node-local Operating System (OS)
13 name of a particular device.

CHAPTER 13

Job Management and Reporting

The job management APIs provide an application with the ability to orchestrate its operation in partnership with the SMS. Members of this category include the [PMIx_Allocation_request](#), [PMIx_Job_control](#), and [PMIx_Process_monitor](#) APIs.

13.1 Allocation Requests

This section defines functionality to request new allocations from the RM, and request modifications to existing allocations. These are primarily used in the following scenarios:

- *Evolving* applications that dynamically request and return resources as they execute.
- *Malleable* environments where the scheduler redirects resources away from executing applications for higher priority jobs or load balancing.
- *Resilient* applications that need to request replacement resources in the face of failures.
- *Rigid* jobs where the user has requested a static allocation of resources for a fixed period of time, but realizes that they underestimated their required time while executing.

PMIx attempts to address this range of use-cases with a flexible API.

13.1.1 PMIx_Allocation_request

Summary

Request an allocation operation from the host resource manager.

Format

PMIx v3.0

C

```
pmix_status_t
PMIx_Allocation_request(pmix_alloc_directive_t directive,
                       pmix_info_t info[], size_t ninfo,
                       pmix_info_t *results[], size_t *nresults);
```

1 **IN directive**
 2 Allocation directive ([pmix_alloc_directive_t](#))
 3 **IN info**
 4 Array of [pmix_info_t](#) structures (array of handles)
 5 **IN ninfo**
 6 Number of elements in the *info* array (integer)
 7 **INOUT results**
 8 Address where a pointer to an array of [pmix_info_t](#) containing the results of the request
 9 can be returned (memory reference)
 10 **INOUT nresults**
 11 Address where the number of elements in *results* can be returned (handle)

Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Required Attributes

13 PMIx libraries are not required to directly support any attributes for this function. However, any
 14 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
 15 *required* to add the [PMIX_USERID](#) and the [PMIX_GRPID](#) attributes of the client process making
 16 the request.

17 Host environments that implement support for this operation are required to support the following
 18 attributes:

19 **[PMIX_ALLOC_REQ_ID](#) "pmix.alloc.reqid" (char*)**
 20 User-provided string identifier for this allocation request which can later be used to query
 21 status of the request.

22 **[PMIX_ALLOC_NUM_NODES](#) "pmix.alloc.nnodes" (uint64_t)**
 23 The number of nodes being requested in an allocation request.

24 **[PMIX_ALLOC_NUM_CPUS](#) "pmix.alloc.ncpus" (uint64_t)**
 25 Number of PUs being requested in an allocation request.

26 **[PMIX_ALLOC_TIME](#) "pmix.alloc.time" (uint32_t)**
 27 Total session time (in seconds) being requested in an allocation request.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_ALLOC_NODE_LIST "pmix.alloc.nlist" (char*)

Regular expression of the specific nodes being requested in an allocation request.

PMIX_ALLOC_NUM_CPU_LIST "pmix.alloc.ncpulist" (char*)

Regular expression of the number of PUs for each node being requested in an allocation request.

PMIX_ALLOC_CPU_LIST "pmix.alloc.cpulist" (char*)

Regular expression of the specific PUs being requested in an allocation request.

PMIX_ALLOC_MEM_SIZE "pmix.alloc.msize" (float)

Number of Megabytes[base2] of memory (per process) being requested in an allocation request.

PMIX_ALLOC_FABRIC "pmix.alloc.net" (array)

Array of `pmix_info_t` describing requested fabric resources. This must include at least:

PMIX_ALLOC_FABRIC_ID, **PMIX_ALLOC_FABRIC_TYPE**, and

PMIX_ALLOC_FABRIC_ENDPTS, plus whatever other descriptors are desired.

PMIX_ALLOC_FABRIC_ID "pmix.alloc.netid" (char*)

The key to be used when accessing this requested fabric allocation. The fabric allocation will be returned/stored as a `pmix_data_array_t` of `pmix_info_t` whose first element is composed of this key and the allocated resource description. The type of the included value depends upon the fabric support. For example, a Transmission Control Protocol (TCP) allocation might consist of a comma-delimited string of socket ranges such as "32000-32100,33005,38123-38146". Additional array entries will consist of any provided resource request directives, along with their assigned values. Examples include: **PMIX_ALLOC_FABRIC_TYPE** - the type of resources provided; **PMIX_ALLOC_FABRIC_PLANE** - if applicable, what plane the resources were assigned from; **PMIX_ALLOC_FABRIC_QOS** - the assigned QoS; **PMIX_ALLOC_BANDWIDTH** - the allocated bandwidth; **PMIX_ALLOC_FABRIC_SEC_KEY** - a security key for the requested fabric allocation. NOTE: the array contents may differ from those requested, especially if **PMIX_INFO_REQD** was not set in the request.

PMIX_ALLOC_BANDWIDTH "pmix.alloc.bw" (float)

Fabric bandwidth (in Megabits[base2]/sec) for the job being requested in an allocation request.

PMIX_ALLOC_FABRIC_QOS "pmix.alloc.netqos" (char*)

Fabric quality of service level for the job being requested in an allocation request.

PMIX_ALLOC_FABRIC_TYPE "pmix.alloc.nettype" (char*)

Type of desired transport (e.g., "tcp", "udp") being requested in an allocation request.

PMIX_ALLOC_FABRIC_PLANE "pmix.alloc.netplane" (char*)

1 ID string for the *fabric plane* to be used for the requested allocation.

2 **PMIX_ALLOC_FABRIC_ENDPTS** "pmix.alloc.endpts" (size_t)

3 Number of endpoints to allocate per *process* in the job.

4 **PMIX_ALLOC_FABRIC_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)

5 Number of endpoints to allocate per *node* for the job.

6 **PMIX_ALLOC_FABRIC_SEC_KEY** "pmix.alloc.nsec" (pmix_byte_object_t)

7 Request that the allocation include a fabric security key for the spawned job.



8 **Description**

9 Request an allocation operation from the host resource manager. Several broad categories are
10 envisioned, including the ability to:

- 11 • Request allocation of additional resources, including memory, bandwidth, and compute. This
12 should be accomplished in a non-blocking manner so that the application can continue to
13 progress while waiting for resources to become available. Note that the new allocation will be
14 disjoint from (i.e., not affiliated with) the allocation of the requestor - thus the termination of one
15 allocation will not impact the other.
- 16 • Extend the reservation on currently allocated resources, subject to scheduling availability and
17 priorities. This includes extending the time limit on current resources, and/or requesting
18 additional resources be allocated to the requesting job. Any additional allocated resources will be
19 considered as part of the current allocation, and thus will be released at the same time.
- 20 • Return no-longer-required resources to the scheduler. This includes the “loan” of resources back
21 to the scheduler with a promise to return them upon subsequent request.

22 If successful, the returned results for a request for additional resources must include the host
23 resource manager’s identifier (**PMIX_ALLOC_ID**) that the requester can use to specify the
24 resources in, for example, a call to **PMIx_Spawn**.

25 **13.1.2 PMIx_Allocation_request_nb**

26 **Summary**

27 Request an allocation operation from the host resource manager.

Format

C

```
pmix_status_t
PMIx_Allocation_request_nb(pmix_alloc_directive_t directive,
                           pmix_info_t info[], size_t ninfo,
                           pmix_info_cbfunc_t cbfunc, void *cbdata);
```

C

IN directive

Allocation directive ([pmix_alloc_directive_t](#))

IN info

Array of [pmix_info_t](#) structures (array of handles)

IN ninfo

Number of elements in the *info* array (integer)

IN cbfunc

Callback function [pmix_info_cbfunc_t](#) (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns one of the following:

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when [PMIX_SUCCESS](#) is returned.

Returns [PMIX_SUCCESS](#) or one of the following error codes when the condition described occurs:

- [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section [3.1.1](#).

Required Attributes

PMIx libraries are not required to directly support any attributes for this function. However, any provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is *required* to add the [PMIX_USERID](#) and the [PMIX_GRPID](#) attributes of the client process making the request.

Host environments that implement support for this operation are required to support the following attributes:

[PMIX_ALLOC_REQ_ID](#) "pmix.alloc.reqid" ([char*](#))

User-provided string identifier for this allocation request which can later be used to query status of the request.

[PMIX_ALLOC_NUM_NODES](#) "pmix.alloc.nnodes" ([uint64_t](#))

1 The number of nodes being requested in an allocation request.
2 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)
3 Number of PUs being requested in an allocation request.
4 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)
5 Total session time (in seconds) being requested in an allocation request.

▲-----▲
▼-----▼ Optional Attributes -----▼

6 The following attributes are optional for host environments that support this operation:
7 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)
8 Regular expression of the specific nodes being requested in an allocation request.
9 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)
10 Regular expression of the number of PUs for each node being requested in an allocation
11 request.
12 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)
13 Regular expression of the specific PUs being requested in an allocation request.
14 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
15 Number of Megabytes[base2] of memory (per process) being requested in an allocation
16 request.
17 **PMIX_ALLOC_FABRIC** "pmix.alloc.net" (array)
18 Array of **pmix_info_t** describing requested fabric resources. This must include at least:
19 **PMIX_ALLOC_FABRIC_ID**, **PMIX_ALLOC_FABRIC_TYPE**, and
20 **PMIX_ALLOC_FABRIC_ENDPTS**, plus whatever other descriptors are desired.
21 **PMIX_ALLOC_FABRIC_ID** "pmix.alloc.netid" (char*)
22 The key to be used when accessing this requested fabric allocation. The fabric allocation
23 will be returned/stored as a **pmix_data_array_t** of **pmix_info_t** whose first
24 element is composed of this key and the allocated resource description. The type of the
25 included value depends upon the fabric support. For example, a TCP allocation might
26 consist of a comma-delimited string of socket ranges such as "**32000-32100,**
27 **33005,38123-38146**". Additional array entries will consist of any provided resource
28 request directives, along with their assigned values. Examples include:
29 **PMIX_ALLOC_FABRIC_TYPE** - the type of resources provided;
30 **PMIX_ALLOC_FABRIC_PLANE** - if applicable, what plane the resources were assigned
31 from; **PMIX_ALLOC_FABRIC_QOS** - the assigned QoS; **PMIX_ALLOC_BANDWIDTH** -
32 the allocated bandwidth; **PMIX_ALLOC_FABRIC_SEC_KEY** - a security key for the
33 requested fabric allocation. NOTE: the array contents may differ from those requested,
34 especially if **PMIX_INFO_REQD** was not set in the request.
35 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)

1 Fabric bandwidth (in Megabits[base2]/sec) for the job being requested in an allocation
2 request.

3 **PMIX_ALLOC_FABRIC_QOS** "pmix.alloc.netqos" (char*)

4 Fabric quality of service level for the job being requested in an allocation request.

5 **PMIX_ALLOC_FABRIC_TYPE** "pmix.alloc.nettype" (char*)

6 Type of desired transport (e.g., "tcp", "udp") being requested in an allocation request.

7 **PMIX_ALLOC_FABRIC_PLANE** "pmix.alloc.netplane" (char*)

8 ID string for the *fabric plane* to be used for the requested allocation.

9 **PMIX_ALLOC_FABRIC_ENDPTS** "pmix.alloc.endpts" (size_t)

10 Number of endpoints to allocate per *process* in the job.

11 **PMIX_ALLOC_FABRIC_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)

12 Number of endpoints to allocate per *node* for the job.

13 **PMIX_ALLOC_FABRIC_SEC_KEY** "pmix.alloc.nsec" (pmix_byte_object_t)

14 Request that the allocation include a fabric security key for the spawned job.



15 Description

16 Non-blocking form of the [PMIx_Allocation_request](#) API.

17 13.1.3 Job Allocation attributes

18 Attributes used to describe the job allocation - these are values passed to and/or returned by the
19 [PMIx_Allocation_request_nb](#) and [PMIx_Allocation_request](#) APIs and are not
20 accessed using the [PMIx_Get](#) API.

21 **PMIX_ALLOC_REQ_ID** "pmix.alloc.reqid" (char*)

22 User-provided string identifier for this allocation request which can later be used to query
23 status of the request.

24 **PMIX_ALLOC_ID** "pmix.alloc.id" (char*)

25 A string identifier (provided by the host environment) for the resulting allocation which can
26 later be used to reference the allocated resources in, for example, a call to [PMIx_Spawn](#).

27 **PMIX_ALLOC_QUEUE** "pmix.alloc.queue" (char*)

28 Name of the WLM queue to which the allocation request is to be directed, or the queue being
29 referenced in a query.

30 **PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)

31 The number of nodes being requested in an allocation request.

32 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)

33 Regular expression of the specific nodes being requested in an allocation request.

34 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)

35 Number of PUs being requested in an allocation request.

1 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)
 2 Regular expression of the number of PUs for each node being requested in an allocation
 3 request.

4 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)
 5 Regular expression of the specific PUs being requested in an allocation request.

6 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
 7 Number of Megabytes[base2] of memory (per process) being requested in an allocation
 8 request.

9 **PMIX_ALLOC_FABRIC** "pmix.alloc.net" (array)
 10 Array of [pmix_info_t](#) describing requested fabric resources. This must include at least:
 11 [PMIX_ALLOC_FABRIC_ID](#), [PMIX_ALLOC_FABRIC_TYPE](#), and
 12 [PMIX_ALLOC_FABRIC_ENDPTS](#), plus whatever other descriptors are desired.

13 **PMIX_ALLOC_FABRIC_ID** "pmix.alloc.netid" (char*)
 14 The key to be used when accessing this requested fabric allocation. The fabric allocation
 15 will be returned/stored as a [pmix_data_array_t](#) of [pmix_info_t](#) whose first
 16 element is composed of this key and the allocated resource description. The type of the
 17 included value depends upon the fabric support. For example, a TCP allocation might
 18 consist of a comma-delimited string of socket ranges such as "32000–32100,
 19 33005, 38123–38146". Additional array entries will consist of any provided resource
 20 request directives, along with their assigned values. Examples include:
 21 [PMIX_ALLOC_FABRIC_TYPE](#) - the type of resources provided;
 22 [PMIX_ALLOC_FABRIC_PLANE](#) - if applicable, what plane the resources were assigned
 23 from; [PMIX_ALLOC_FABRIC_QOS](#) - the assigned QoS; [PMIX_ALLOC_BANDWIDTH](#) -
 24 the allocated bandwidth; [PMIX_ALLOC_FABRIC_SEC_KEY](#) - a security key for the
 25 requested fabric allocation. NOTE: the array contents may differ from those requested,
 26 especially if [PMIX_INFO_REQD](#) was not set in the request.

27 **PMIX_ALLOC_BANDWIDTH** "pmix.alloc.bw" (float)
 28 Fabric bandwidth (in Megabits[base2]/sec) for the job being requested in an allocation
 29 request.

30 **PMIX_ALLOC_FABRIC_QOS** "pmix.alloc.netqos" (char*)
 31 Fabric quality of service level for the job being requested in an allocation request.

32 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)
 33 Total session time (in seconds) being requested in an allocation request.

34 **PMIX_ALLOC_FABRIC_TYPE** "pmix.alloc.nettype" (char*)
 35 Type of desired transport (e.g., "tcp", "udp") being requested in an allocation request.

36 **PMIX_ALLOC_FABRIC_PLANE** "pmix.alloc.netplane" (char*)
 37 ID string for the *fabric plane* to be used for the requested allocation.

38 **PMIX_ALLOC_FABRIC_ENDPTS** "pmix.alloc.endpts" (size_t)
 39 Number of endpoints to allocate per *process* in the job.

40 **PMIX_ALLOC_FABRIC_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)
 41 Number of endpoints to allocate per *node* for the job.

42 **PMIX_ALLOC_FABRIC_SEC_KEY** "pmix.alloc.nsec" (pmix_byte_object_t)
 43 Request that the allocation include a fabric security key for the spawned job.

1 13.1.4 Job Allocation Directives

2 The `pmix_alloc_directive_t` structure is a `uint8_t` type that defines the behavior of
3 allocation requests. The following constants can be used to set a variable of the type
4 `pmix_alloc_directive_t`. All definitions were introduced in version 2 of the standard
5 unless otherwise marked.

6 **PMIX_ALLOC_NEW 1** A new allocation is being requested. The resulting allocation will be
7 disjoint (i.e., not connected in a job sense) from the requesting allocation.

8 **PMIX_ALLOC_EXTEND 2** Extend the existing allocation, either in time or as additional
9 resources.

10 **PMIX_ALLOC_RELEASE 3** Release part of the existing allocation. Attributes in the
11 accompanying `pmix_info_t` array may be used to specify permanent release of the
12 identified resources, or “lending” of those resources for some period of time.

13 **PMIX_ALLOC_REAQUIRE 4** Reacquire resources that were previously “lent” back to the
14 scheduler.

15 **PMIX_ALLOC_EXTERNAL** A value boundary above which implementers are free to define
16 their own directive values.

17 13.2 Job Control

18 This section defines APIs that enable the application and host environment to coordinate the
19 response to failures and other events. This can include requesting termination of the entire job or a
20 subset of processes within a job, but can also be used in combination with other PMIx capabilities
21 (e.g., allocation support and event notification) for more nuanced responses. For example, an
22 application notified of an incipient over-temperature condition on a node could use the
23 `PMIx_Allocation_request_nb` interface to request replacement nodes while
24 simultaneously using the `PMIx_Job_control_nb` interface to direct that a checkpoint event be
25 delivered to all processes in the application. If replacement resources are not available, the
26 application might use the `PMIx_Job_control_nb` interface to request that the job continue at a
27 lower power setting, perhaps sufficient to avoid the over-temperature failure.

28 The job control APIs can also be used by an application to register itself as available for preemption
29 when operating in an environment such as a cloud or where incentives, financial or otherwise, are
30 provided to jobs willing to be preempted. Registration can include attributes indicating how many
31 resources are being offered for preemption (e.g., all or only some portion), whether the application
32 will require time to prepare for preemption, etc. Jobs that request a warning will receive an event
33 notifying them of an impending preemption (possibly including information as to the resources that
34 will be taken away, how much time the application will be given prior to being preempted, whether
35 the preemption will be a suspension or full termination, etc.) so they have an opportunity to save
36 their work. Once the application is ready, it calls the provided event completion callback function to
37 indicate that the SMS is free to suspend or terminate it, and can include directives regarding any
38 desired restart.

1 13.2.1 PMIx_Job_control

2 Summary

3 Request a job control action.

4 Format

PMIx v3.0

C

```
5 pmix_status_t
6 PMIx_Job_control(const pmix_proc_t targets[], size_t ntargets,
7                 const pmix_info_t directives[], size_t ndirs,
8                 pmix_info_t *results[], size_t *nresults);
```

C

9 IN targets

10 Array of proc structures (array of handles)

11 IN ntargets

12 Number of elements in the *targets* array (integer)

13 IN directives

14 Array of info structures (array of handles)

15 IN ndirs

16 Number of elements in the *directives* array (integer)

17 INOUT results

18 Address where a pointer to an array of `pmix_info_t` containing the results of the request
19 can be returned (memory reference)

20 INOUT nresults

21 Address where the number of elements in *results* can be returned (handle)

22 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

23 PMIx libraries are not required to directly support any attributes for this function. However, any
24 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
25 *required* to add the `PMIX_USERID` and the `PMIX_GRPID` attributes of the client process making
26 the request.

27 Host environments that implement support for this operation are required to support the following
28 attributes:

29 `PMIX_JOB_CTRL_ID` "pmix.jctrl.id" (char*)

30 Provide a string identifier for this request. The user can provide an identifier for the
31 requested operation, thus allowing them to later request status of the operation or to
32 terminate it. The host, therefore, shall track it with the request for future reference.

33 `PMIX_JOB_CTRL_PAUSE` "pmix.jctrl.pause" (bool)

34 Pause the specified processes.

35 `PMIX_JOB_CTRL_RESUME` "pmix.jctrl.resume" (bool)

1 Resume (“un-pause”) the specified processes.

2 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)

3 Forcibly terminate the specified processes and cleanup.

4 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)

5 Send given signal to specified processes.

6 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)

7 Politely terminate the specified processes.

8 **PMIX_REGISTER_CLEANUP** "pmix.reg.cleanup" (char*)

9 Comma-delimited list of files to be removed upon process termination.

10 **PMIX_REGISTER_CLEANUP_DIR** "pmix.reg.cleanupdir" (char*)

11 Comma-delimited list of directories to be removed upon process termination.

12 **PMIX_CLEANUP_RECURSIVE** "pmix.clnup.recurse" (bool)

13 Recursively cleanup all subdirectories under the specified one(s).

14 **PMIX_CLEANUP_EMPTY** "pmix.clnup.empty" (bool)

15 Only remove empty subdirectories.

16 **PMIX_CLEANUP_IGNORE** "pmix.clnup.ignore" (char*)

17 Comma-delimited list of filenames that are not to be removed.

18 **PMIX_CLEANUP_LEAVE_TOPDIR** "pmix.clnup.lvtop" (bool)

19 When recursively cleaning subdirectories, do not remove the top-level directory (the one
20 given in the cleanup request).

▲-----▲
▼-----▼ **Optional Attributes** -----▼

21 The following attributes are optional for host environments that support this operation:

22 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)

23 Cancel the specified request - the provided request ID must match the
24 **PMIX_JOB_CTRL_ID** provided to a previous call to **PMIx_Job_control**. An ID of
25 **NULL** implies cancel all requests from this requestor.

26 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)

27 Restart the specified processes using the given checkpoint ID.

28 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)

29 Checkpoint the specified processes and assign the given ID to it.

30 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)

31 Use event notification to trigger a process checkpoint.

32 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)

33 Use the given signal to trigger a process checkpoint.

1 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)

2 Time in seconds to wait for a checkpoint to complete.

3 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**

4 "pmix.jctrl.ckmethod" (pmix_data_array_t)

5 Array of **pmix_info_t** declaring each method and value supported by this application.

6 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)

7 Regular expression identifying nodes that are to be provisioned.

8 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)

9 Name of the image that is to be provisioned.

10 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)

11 Indicate that the job can be pre-empted.



12 Description

13 Request a job control action. The *targets* array identifies the processes to which the requested job
14 control action is to be applied. All *clones* of an identified process are to have the requested action
15 applied to them. A **NULL** value can be used to indicate all processes in the caller's namespace. The
16 use of **PMIX_RANK_WILDCARD** can also be used to indicate that all processes in the given
17 namespace are to be included.

18 The directives are provided as **pmix_info_t** structures in the *directives* array. The returned
19 *status* indicates whether or not the request was granted, and information as to the reason for any
20 denial of the request shall be returned in the *results* array.

21 13.2.2 PMIx_Job_control_nb

22 Summary

23 Request a job control action.

24 Format

PMIx v2.0

C

25 **pmix_status_t**

26 **PMIx_Job_control_nb**(const pmix_proc_t targets[], size_t ntargets,
27 const pmix_info_t directives[], size_t ndirs,
28 pmix_info_cbfunc_t cbfunc, void *cbdata);

1 **IN targets**
 2 Array of proc structures (array of handles)
 3 **IN ntargets**
 4 Number of elements in the *targets* array (integer)
 5 **IN directives**
 6 Array of info structures (array of handles)
 7 **IN ndirs**
 8 Number of elements in the *directives* array (integer)
 9 **IN cbfunc**
 10 Callback function `pmix_info_cbfunc_t` (function reference)
 11 **IN cbdata**
 12 Data to be passed to the callback function (memory reference)

13 A successful return indicates that the request is being processed and the result will be returned in
 14 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
 15 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

16 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 17 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
 18 returned *success* - the *cbfunc* will *not* be called

19 If none of the above return codes are appropriate, then an implementation must return either a
 20 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

21 PMIx libraries are not required to directly support any attributes for this function. However, any
 22 provided attributes must be passed to the host SMS daemon for processing, and the PMIx library is
 23 *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the client process making
 24 the request.

25 Host environments that implement support for this operation are required to support the following
 26 attributes:

27 **PMIX_JOB_CTRL_ID** "pmix.jctrl.id" (char*)

28 Provide a string identifier for this request. The user can provide an identifier for the
 29 requested operation, thus allowing them to later request status of the operation or to
 30 terminate it. The host, therefore, shall track it with the request for future reference.

31 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)

32 Pause the specified processes.

33 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)

34 Resume ("un-pause") the specified processes.

35 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)

1 Forcibly terminate the specified processes and cleanup.

2 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)

3 Send given signal to specified processes.

4 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)

5 Politely terminate the specified processes.

6 **PMIX_REGISTER_CLEANUP** "pmix.reg.cleanup" (char*)

7 Comma-delimited list of files to be removed upon process termination.

8 **PMIX_REGISTER_CLEANUP_DIR** "pmix.reg.cleanupdir" (char*)

9 Comma-delimited list of directories to be removed upon process termination.

10 **PMIX_CLEANUP_RECURSIVE** "pmix.clnup.recurse" (bool)

11 Recursively cleanup all subdirectories under the specified one(s).

12 **PMIX_CLEANUP_EMPTY** "pmix.clnup.empty" (bool)

13 Only remove empty subdirectories.

14 **PMIX_CLEANUP_IGNORE** "pmix.clnup.ignore" (char*)

15 Comma-delimited list of filenames that are not to be removed.

16 **PMIX_CLEANUP_LEAVE_TOPDIR** "pmix.clnup.lvtop" (bool)

17 When recursively cleaning subdirectories, do not remove the top-level directory (the one
18 given in the cleanup request).

▲-----▲
▼-----▼ **Optional Attributes** -----▼

19 The following attributes are optional for host environments that support this operation:

20 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)

21 Cancel the specified request - the provided request ID must match the

22 **PMIX_JOB_CTRL_ID** provided to a previous call to **PMIx_Job_control**. An ID of
23 **NULL** implies cancel all requests from this requestor.

24 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)

25 Restart the specified processes using the given checkpoint ID.

26 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)

27 Checkpoint the specified processes and assign the given ID to it.

28 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)

29 Use event notification to trigger a process checkpoint.

30 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)

31 Use the given signal to trigger a process checkpoint.

32 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)

33 Time in seconds to wait for a checkpoint to complete.

1 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
 2 "pmix.jctrl.ckmethod" (pmix_data_array_t)
 3 Array of pmix_info_t declaring each method and value supported by this application.
 4 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
 5 Regular expression identifying nodes that are to be provisioned.
 6 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)
 7 Name of the image that is to be provisioned.
 8 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)
 9 Indicate that the job can be pre-empted.



Description

Non-blocking form of the **PMIx_Job_control** API. The *targets* array identifies the processes to which the requested job control action is to be applied. All *clones* of an identified process are to have the requested action applied to them. A **NULL** value can be used to indicate all processes in the caller's namespace. The use of **PMIX_RANK_WILDCARD** can also be used to indicate that all processes in the given namespace are to be included.

The directives are provided as **pmix_info_t** structures in the *directives* array. The callback function provides a *status* to indicate whether or not the request was granted, and to provide some information as to the reason for any denial in the **pmix_info_cbfnc_t** array of **pmix_info_t** structures.

13.2.3 Job control constants

The following constants are specifically defined for return by the job control APIs:

PMIX_ERR_CONFLICTING_CLEANUP_DIRECTIVES -51 Conflicting directives given for job/process cleanup.

13.2.4 Job control events

The following job control events may be available for registration, depending upon implementation and host environment support:

PMIX_JCTRL_CHECKPOINT -106 Monitored by PMIx client to trigger a checkpoint operation.
PMIX_JCTRL_CHECKPOINT_COMPLETE -107 Sent by a PMIx client and monitored by a PMIx server to notify that requested checkpoint operation has completed.
PMIX_JCTRL_PREEMPT_ALERT -108 Monitored by a PMIx client to detect that an RM intends to preempt the job.
PMIX_ERR_PROC_RESTART -4 Error in process restart.
PMIX_ERR_PROC_CHECKPOINT -5 Error in process checkpoint.
PMIX_ERR_PROC_MIGRATE -6 Error in process migration.

1 13.2.5 Job control attributes

2 Attributes used to request control operations on an executing application - these are values passed
3 to the job control APIs and are not accessed using the [PMIx_Get](#) API.

4 **PMIX_JOB_CTRL_ID** "pmix.jctrl.id" (char*)
5 Provide a string identifier for this request. The user can provide an identifier for the
6 requested operation, thus allowing them to later request status of the operation or to
7 terminate it. The host, therefore, shall track it with the request for future reference.

8 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)
9 Pause the specified processes.

10 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)
11 Resume ("un-pause") the specified processes.

12 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)
13 Cancel the specified request - the provided request ID must match the
14 [PMIX_JOB_CTRL_ID](#) provided to a previous call to [PMIx_Job_control](#). An ID of
15 NULL implies cancel all requests from this requestor.

16 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)
17 Forcibly terminate the specified processes and cleanup.

18 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)
19 Restart the specified processes using the given checkpoint ID.

20 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)
21 Checkpoint the specified processes and assign the given ID to it.

22 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)
23 Use event notification to trigger a process checkpoint.

24 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)
25 Use the given signal to trigger a process checkpoint.

26 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)
27 Time in seconds to wait for a checkpoint to complete.

28 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
29 "pmix.jctrl.ckmethod" (pmix_data_array_t)
30 Array of [pmix_info_t](#) declaring each method and value supported by this application.

31 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)
32 Send given signal to specified processes.

33 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
34 Regular expression identifying nodes that are to be provisioned.

35 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvnimg" (char*)
36 Name of the image that is to be provisioned.

37 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)
38 Indicate that the job can be pre-empted.

39 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)
40 Politely terminate the specified processes.

41 **PMIX_REGISTER_CLEANUP** "pmix.reg.cleanup" (char*)
42 Comma-delimited list of files to be removed upon process termination.

```

1  PMIX_REGISTER_CLEANUP_DIR "pmix.reg.cleanupdir" (char*)
2      Comma-delimited list of directories to be removed upon process termination.
3  PMIX_CLEANUP_RECURSIVE "pmix.clnup.recurse" (bool)
4      Recursively cleanup all subdirectories under the specified one(s).
5  PMIX_CLEANUP_EMPTY "pmix.clnup.empty" (bool)
6      Only remove empty subdirectories.
7  PMIX_CLEANUP_IGNORE "pmix.clnup.ignore" (char*)
8      Comma-delimited list of filenames that are not to be removed.
9  PMIX_CLEANUP_LEAVE_TOPDIR "pmix.clnup.lvtop" (bool)
10     When recursively cleaning subdirectories, do not remove the top-level directory (the one
11     given in the cleanup request).

```

12 13.3 Process and Job Monitoring

13 In addition to external faults, a common problem encountered in HPC applications is a failure to
14 make progress due to some internal conflict in the computation. These situations can result in a
15 significant waste of resources as the SMS is unaware of the problem, and thus cannot terminate the
16 job. Various watchdog methods have been developed for detecting this situation, including
17 requiring a periodic “heartbeat” from the application and monitoring a specified file for changes in
18 size and/or modification time.

19 The following APIs allow applications to request monitoring, directing what is to be monitored, the
20 frequency of the associated check, whether or not the application is to be notified (via the event
21 notification subsystem) of stall detection, and other characteristics of the operation.

22 13.3.1 PMIx_Process_monitor

23 Summary

24 Request that application processes be monitored.

25 *PMIx v3.0* Format

```

26 pmix_status_t
27 PMIx_Process_monitor(const pmix_info_t *monitor,
28                     pmix_status_t error,
29                     const pmix_info_t directives[], size_t ndirs,
30                     pmix_info_t *results[], size_t *nresults);

```

```

31 IN monitor
32     info (handle)
33 IN error
34     status (integer)

```

1 **IN directives**

2 Array of info structures (array of handles)

3 **IN ndirs**

4 Number of elements in the *directives* array (integer)

5 **INOUT results**

6 Address where a pointer to an array of `pmix_info_t` containing the results of the request
7 can be returned (memory reference)

8 **INOUT nresults**

9 Address where the number of elements in *results* can be returned (handle)

10 A successful return indicates that the results have been placed in the *results* array.

11 Returns `PMIX_SUCCESS` or a negative value indicating the error.

▼----- Optional Attributes -----▼

12 The following attributes may be implemented by a PMIx library or by the host environment. If
13 supported by the PMIx server library, then the library must not pass the supported attributes to the
14 host environment. All attributes not directly supported by the server library must be passed to the
15 host environment if it supports this operation, and the library is *required* to add the
16 `PMIX_USERID` and the `PMIX_GRPID` attributes of the requesting process:

17 `PMIX_MONITOR_ID` "pmix.monitor.id" (char*)

18 Provide a string identifier for this request.

19 `PMIX_MONITOR_CANCEL` "pmix.monitor.cancel" (char*)

20 Identifier to be canceled (`NULL` means cancel all monitoring for this process).

21 `PMIX_MONITOR_APP_CONTROL` "pmix.monitor.appctrl" (bool)

22 The application desires to control the response to a monitoring event - i.e., the application is
23 requesting that the host environment not take immediate action in response to the event (e.g.,
24 terminating the job).

25 `PMIX_MONITOR_HEARTBEAT` "pmix.monitor.mbeat" (void)

26 Register to have the PMIx server monitor the requestor for heartbeats.

27 `PMIX_MONITOR_HEARTBEAT_TIME` "pmix.monitor.btime" (uint32_t)

28 Time in seconds before declaring heartbeat missed.

29 `PMIX_MONITOR_HEARTBEAT_DROPS` "pmix.monitor.bdrop" (uint32_t)

30 Number of heartbeats that can be missed before generating the event.

31 `PMIX_MONITOR_FILE` "pmix.monitor.fmon" (char*)

32 Register to monitor file for signs of life.

33 `PMIX_MONITOR_FILE_SIZE` "pmix.monitor.fsize" (bool)

34 Monitor size of given file is growing to determine if the application is running.

35 `PMIX_MONITOR_FILE_ACCESS` "pmix.monitor.faccess" (char*)

36 Monitor time since last access of given file to determine if the application is running.

1 **PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)
2 Monitor time since last modified of given file to determine if the application is running.
3 **PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)
4 Time in seconds between checking the file.
5 **PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrop" (uint32_t)
6 Number of file checks that can be missed before generating the event.
7 **PMIX_SEND_HEARTBEAT** "pmix.monitor.beat" (void)
8 Send heartbeat to local PMIx server.

Description

Request that application processes be monitored via several possible methods. For example, that the server monitor this process for periodic heartbeats as an indication that the process has not become “wedged”. When a monitor detects the specified alarm condition, it will generate an event notification using the provided error code and passing along any available relevant information. It is up to the caller to register a corresponding event handler.

The *monitor* argument is an attribute indicating the type of monitor being requested. For example, **PMIX_MONITOR_FILE** to indicate that the requestor is asking that a file be monitored.

The *error* argument is the status code to be used when generating an event notification alerting that the monitor has been triggered. The range of the notification defaults to **PMIX_RANGE_NAMESPACE**. This can be changed by providing a **PMIX_RANGE** directive.

The *directives* argument characterizes the monitoring request (e.g., monitor file size) and frequency of checking to be done

The returned *status* indicates whether or not the request was granted, and information as to the reason for any denial of the request shall be returned in the *results* array.

13.3.2 PMIx_Process_monitor_nb

Summary

Request that application processes be monitored.

Format

```
pmix_status_t
PMIx_Process_monitor_nb(const pmix_info_t *monitor,
                        pmix_status_t error,
                        const pmix_info_t directives[],
                        size_t ndirs,
                        pmix_info_cbfunc_t cbfunc, void *cbdata);
```

- IN monitor**
info (handle)
- IN error**
status (integer)
- IN directives**
Array of info structures (array of handles)
- IN ndirs**
Number of elements in the *directives* array (integer)
- IN cbfunc**
Callback function [pmix_info_cbfunc_t](#) (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

Returns one of the following:

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will *not* be called

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Optional Attributes

The following attributes may be implemented by a PMIx library or by the host environment. If supported by the PMIx server library, then the library must not pass the supported attributes to the host environment. All attributes not directly supported by the server library must be passed to the host environment if it supports this operation, and the library is *required* to add the **PMIX_USERID** and the **PMIX_GRPID** attributes of the requesting process:

PMIX_MONITOR_ID "pmix.monitor.id" (char*)
Provide a string identifier for this request.

1 **PMIX_MONITOR_CANCEL** "pmix.monitor.cancel" (char*)
 2 Identifier to be canceled (NULL means cancel all monitoring for this process).

3 **PMIX_MONITOR_APP_CONTROL** "pmix.monitor.appctrl" (bool)
 4 The application desires to control the response to a monitoring event - i.e., the application is
 5 requesting that the host environment not take immediate action in response to the event (e.g.,
 6 terminating the job).

7 **PMIX_MONITOR_HEARTBEAT** "pmix.monitor.mbeat" (void)
 8 Register to have the PMIx server monitor the requestor for heartbeats.

9 **PMIX_MONITOR_HEARTBEAT_TIME** "pmix.monitor.btime" (uint32_t)
 10 Time in seconds before declaring heartbeat missed.

11 **PMIX_MONITOR_HEARTBEAT_DROPS** "pmix.monitor.bdrop" (uint32_t)
 12 Number of heartbeats that can be missed before generating the event.

13 **PMIX_MONITOR_FILE** "pmix.monitor.fmon" (char*)
 14 Register to monitor file for signs of life.

15 **PMIX_MONITOR_FILE_SIZE** "pmix.monitor.fsize" (bool)
 16 Monitor size of given file is growing to determine if the application is running.

17 **PMIX_MONITOR_FILE_ACCESS** "pmix.monitor.faccess" (char*)
 18 Monitor time since last access of given file to determine if the application is running.

19 **PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)
 20 Monitor time since last modified of given file to determine if the application is running.

21 **PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)
 22 Time in seconds between checking the file.

23 **PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrop" (uint32_t)
 24 Number of file checks that can be missed before generating the event.

25 **PMIX_SEND_HEARTBEAT** "pmix.monitor.beat" (void)
 26 Send heartbeat to local PMIx server.

▲-----▲

27 Description

28 Non-blocking form of the **PMIx_Process_monitor** API. The *cbfunc* function provides a
 29 *status* to indicate whether or not the request was granted, and to provide some information as to the
 30 reason for any denial in the **pmix_info_cbfunc_t** array of **pmix_info_t** structures.

31 13.3.3 PMIx_Heartbeat

32 Summary

33 Send a heartbeat to the PMIx server library

1 **Format** C

2 `PMIx_Heartbeat ();` C

3 **Description**

4 A simplified macro wrapping `PMIx_Process_monitor_nb` that sends a heartbeat to the PMIx
5 server library.

6 **13.3.4 Monitoring events**

7 The following monitoring events may be available for registration, depending upon implementation
8 and host environment support:

- 9 **PMIX_MONITOR_HEARTBEAT_ALERT -109** Heartbeat failed to arrive within specified
10 window. The process that triggered this alert will be identified in the event.
11 **PMIX_MONITOR_FILE_ALERT -110** File failed its monitoring detection criteria. The file
12 that triggered this alert will be identified in the event.

13 **13.3.5 Monitoring attributes**

14 Attributes used to control monitoring of an executing application- these are values passed to the
15 `PMIx_Process_monitor_nb` API and are not accessed using the `PMIx_Get` API.

- 16 **PMIX_MONITOR_ID "pmix.monitor.id" (char*)**
17 Provide a string identifier for this request.
- 18 **PMIX_MONITOR_CANCEL "pmix.monitor.cancel" (char*)**
19 Identifier to be canceled (`NULL` means cancel all monitoring for this process).
- 20 **PMIX_MONITOR_APP_CONTROL "pmix.monitor.appctrl" (bool)**
21 The application desires to control the response to a monitoring event - i.e., the application is
22 requesting that the host environment not take immediate action in response to the event (e.g.,
23 terminating the job).
- 24 **PMIX_MONITOR_HEARTBEAT "pmix.monitor.mbeat" (void)**
25 Register to have the PMIx server monitor the requestor for heartbeats.
- 26 **PMIX_SEND_HEARTBEAT "pmix.monitor.beat" (void)**
27 Send heartbeat to local PMIx server.
- 28 **PMIX_MONITOR_HEARTBEAT_TIME "pmix.monitor.btime" (uint32_t)**
29 Time in seconds before declaring heartbeat missed.
- 30 **PMIX_MONITOR_HEARTBEAT_DROPS "pmix.monitor.bdrop" (uint32_t)**
31 Number of heartbeats that can be missed before generating the event.
- 32 **PMIX_MONITOR_FILE "pmix.monitor.fmon" (char*)**
33 Register to monitor file for signs of life.
- 34 **PMIX_MONITOR_FILE_SIZE "pmix.monitor.fsize" (bool)**
35 Monitor size of given file is growing to determine if the application is running.

```

1 PMIX_MONITOR_FILE_ACCESS "pmix.monitor.faccess" (char*)
2     Monitor time since last access of given file to determine if the application is running.
3 PMIX_MONITOR_FILE_MODIFY "pmix.monitor.fmod" (char*)
4     Monitor time since last modified of given file to determine if the application is running.
5 PMIX_MONITOR_FILE_CHECK_TIME "pmix.monitor.ftime" (uint32_t)
6     Time in seconds between checking the file.
7 PMIX_MONITOR_FILE_DROPS "pmix.monitor.fdrop" (uint32_t)
8     Number of file checks that can be missed before generating the event.

```

9 13.4 Logging

10 The logging interface supports posting information by applications and SMS elements to persistent
11 storage. This function is *not* intended for output of computational results, but rather for reporting
12 status and saving state information such as inserting computation progress reports into the
13 application's SMS job log or error reports to the local syslog.

14 13.4.1 PMIx_Log

15 Summary

16 Log data to a data service.

17 Format

PMIx v3.0

```

18 pmix_status_t
19 PMIx_Log(const pmix_info_t data[], size_t ndata,
20         const pmix_info_t directives[], size_t ndirs);

```

21 **IN data**
22 Array of info structures (array of handles)

23 **IN ndata**
24 Number of elements in the *data* array (**size_t**)

25 **IN directives**
26 Array of info structures (array of handles)

27 **IN ndirs**
28 Number of elements in the *directives* array (**size_t**)

29 Returns **PMIX_SUCCESS** or a negative value indicating the error.

Required Attributes

1 If the PMIx library does not itself perform this operation, then it is required to pass any attributes
2 provided by the client to the host environment for processing. In addition, it must include the
3 following attributes in the passed *info* array:

4 **PMIX_USERID** "pmix.euid" (uint32_t)

5 Effective user ID of the connecting process.

6 **PMIX_GRPID** "pmix.egid" (uint32_t)

7 Effective group ID of the connecting process.

8 Host environments or PMIx libraries that implement support for this operation are required to
9 support the following attributes:

10 **PMIX_LOG_STDERR** "pmix.log.stderr" (char*)

11 Log string to **stderr**.

12 **PMIX_LOG_STDOUT** "pmix.log.stdout" (char*)

13 Log string to **stdout**.

14 **PMIX_LOG_SYSLOG** "pmix.log.syslog" (char*)

15 Log data to syslog. Defaults to **ERROR** priority. Will log to global syslog if available,
16 otherwise to local syslog.

17 **PMIX_LOG_LOCAL_SYSLOG** "pmix.log.lsys" (char*)

18 Log data to local syslog. Defaults to **ERROR** priority.

19 **PMIX_LOG_GLOBAL_SYSLOG** "pmix.log.gsys" (char*)

20 Forward data to system "gateway" and log msg to that syslog Defaults to **ERROR** priority.

21 **PMIX_LOG_SYSLOG_PRI** "pmix.log.syspri" (int)

22 Syslog priority level.

23 **PMIX_LOG_ONCE** "pmix.log.once" (bool)

24 Only log this once with whichever channel can first support it, taking the channels in priority
25 order.

Optional Attributes

26 The following attributes are optional for host environments or PMIx libraries that support this
27 operation:

28 **PMIX_LOG_SOURCE** "pmix.log.source" (pmix_proc_t*)

29 ID of source of the log request.

30 **PMIX_LOG_TIMESTAMP** "pmix.log.tstamp" (time_t)

31 Timestamp for log report.

32 **PMIX_LOG_GENERATE_TIMESTAMP** "pmix.log.gtstamp" (bool)

1 Generate timestamp for log.

2 **PMIX_LOG_TAG_OUTPUT** "pmix.log.tag" (bool)

3 Label the output stream with the channel name (e.g., "stdout").

4 **PMIX_LOG_TIMESTAMP_OUTPUT** "pmix.log.tsout" (bool)

5 Print timestamp in output string.

6 **PMIX_LOG_XML_OUTPUT** "pmix.log.xml" (bool)

7 Print the output stream in eXtensible Markup Language (XML) format.

8 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)

9 Log via email based on **pmix_info_t** containing directives.

10 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)

11 Comma-delimited list of email addresses that are to receive the message.

12 **PMIX_LOG_EMAIL_SENDER_ADDR** "pmix.log.emfaddr" (char*)

13 Return email address of sender.

14 **PMIX_LOG_EMAIL_SERVER** "pmix.log.esrvr" (char*)

15 Hostname (or IP address) of SMTP server.

16 **PMIX_LOG_EMAIL_SRVR_PORT** "pmix.log.esrvrprt" (int32_t)

17 Port the email server is listening to.

18 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)

19 Subject line for email.

20 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)

21 Message to be included in email.

22 **PMIX_LOG_JOB_RECORD** "pmix.log.jrec" (bool)

23 Log the provided information to the host environment's job record.

24 **PMIX_LOG_GLOBAL_DATASTORE** "pmix.log.gstore" (bool)

25 Store the log data in a global data store (e.g., database).



26 **Description**

27 Log data subject to the services offered by the host environment. The data to be logged is provided
28 in the *data* array. The (optional) *directives* can be used to direct the choice of logging channel.

29 **Advice to users**

30 It is strongly recommended that the **PMIx_Log** API not be used by applications for streaming data
31 as it is not a "performant" transport and can perturb the application since it involves the local PMIx
32 server and host SMS daemon. Note that a return of **PMIX_SUCCESS** only denotes that the data
33 was successfully handed to the appropriate system call (for local channels) or the host environment
 and does not indicate receipt at the final destination.

1 13.4.2 PMI_x_Log_nb

2 Summary

3 Log data to a data service.

4 Format

PMIx v2.0

C

5 `pmix_status_t`

```
6 PMIx_Log_nb(const pmix_info_t data[], size_t ndata,  
7             const pmix_info_t directives[], size_t ndirs,  
8             pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

9 **IN data**

Array of info structures (array of handles)

11 **IN ndata**

Number of elements in the *data* array (**size_t**)

13 **IN directives**

Array of info structures (array of handles)

15 **IN ndirs**

Number of elements in the *directives* array (**size_t**)

17 **IN cbfunc**

Callback function [pmix_op_cbfunc_t](#) (function reference)

19 **IN cbdata**

Data to be passed to the callback function (memory reference)

21 Return codes are one of the following:

22 A successful return indicates that the request is being processed and the result will be returned in
23 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
24 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

25 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

26 **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
27 returned *success* - the *cbfunc* will *not* be called

28 **PMIX_ERR_BAD_PARAM** The logging request contains at least one incorrect entry that prevents
29 it from being processed. The callback function will not be called.

30 If none of the above return codes are appropriate, then an implementation must return either a
31 general PMI_x error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

1 If the PMIx library does not itself perform this operation, then it is required to pass any attributes
2 provided by the client to the host environment for processing. In addition, it must include the
3 following attributes in the passed *info* array:

4 **PMIX_USERID** "pmix.euid" (uint32_t)

5 Effective user ID of the connecting process.

6 **PMIX_GRPID** "pmix.egid" (uint32_t)

7 Effective group ID of the connecting process.

8 Host environments or PMIx libraries that implement support for this operation are required to
9 support the following attributes:

10 **PMIX_LOG_STDERR** "pmix.log.stderr" (char*)

11 Log string to **stderr**.

12 **PMIX_LOG_STDOUT** "pmix.log.stdout" (char*)

13 Log string to **stdout**.

14 **PMIX_LOG_SYSLOG** "pmix.log.syslog" (char*)

15 Log data to syslog. Defaults to **ERROR** priority. Will log to global syslog if available,
16 otherwise to local syslog.

17 **PMIX_LOG_LOCAL_SYSLOG** "pmix.log.lsys" (char*)

18 Log data to local syslog. Defaults to **ERROR** priority.

19 **PMIX_LOG_GLOBAL_SYSLOG** "pmix.log.gsys" (char*)

20 Forward data to system "gateway" and log msg to that syslog Defaults to **ERROR** priority.

21 **PMIX_LOG_SYSLOG_PRI** "pmix.log.syspri" (int)

22 Syslog priority level.

23 **PMIX_LOG_ONCE** "pmix.log.once" (bool)

24 Only log this once with whichever channel can first support it, taking the channels in priority
25 order.

Optional Attributes

26 The following attributes are optional for host environments or PMIx libraries that support this
27 operation:

28 **PMIX_LOG_SOURCE** "pmix.log.source" (pmix_proc_t*)

29 ID of source of the log request.

30 **PMIX_LOG_TIMESTAMP** "pmix.log.tstamp" (time_t)

31 Timestamp for log report.

32 **PMIX_LOG_GENERATE_TIMESTAMP** "pmix.log.gtstamp" (bool)

1 Generate timestamp for log.

2 **PMIX_LOG_TAG_OUTPUT** "pmix.log.tag" (bool)

3 Label the output stream with the channel name (e.g., "stdout").

4 **PMIX_LOG_TIMESTAMP_OUTPUT** "pmix.log.tsout" (bool)

5 Print timestamp in output string.

6 **PMIX_LOG_XML_OUTPUT** "pmix.log.xml" (bool)

7 Print the output stream in XML format.

8 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)

9 Log via email based on **pmix_info_t** containing directives.

10 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)

11 Comma-delimited list of email addresses that are to receive the message.

12 **PMIX_LOG_EMAIL_SENDER_ADDR** "pmix.log.emfaddr" (char*)

13 Return email address of sender.

14 **PMIX_LOG_EMAIL_SERVER** "pmix.log.esrvr" (char*)

15 Hostname (or IP address) of SMTP server.

16 **PMIX_LOG_EMAIL_SRVR_PORT** "pmix.log.esrvrprt" (int32_t)

17 Port the email server is listening to.

18 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)

19 Subject line for email.

20 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)

21 Message to be included in email.

22 **PMIX_LOG_JOB_RECORD** "pmix.log.jrec" (bool)

23 Log the provided information to the host environment's job record.

24 **PMIX_LOG_GLOBAL_DATASTORE** "pmix.log.gstore" (bool)

25 Store the log data in a global data store (e.g., database).



26 **Description**

27 Log data subject to the services offered by the host environment. The data to be logged is provided
28 in the *data* array. The (optional) *directives* can be used to direct the choice of logging channel. The
29 callback function will be executed when the log operation has been completed. The *data* and
30 *directives* arrays must be maintained until the callback is provided.

Advice to users

It is strongly recommended that the `PMix_Log_nb` API not be used by applications for streaming data as it is not a “performant” transport and can perturb the application since it involves the local PMix server and host SMS daemon. Note that a return of `PMIX_SUCCESS` only denotes that the data was successfully handed to the appropriate system call (for local channels) or the host environment and does not indicate receipt at the final destination.

13.4.3 Log attributes

Attributes used to describe `PMix_Log` behavior - these are values passed to the `PMix_Log` API and therefore are not accessed using the `PMix_Get` API.

`PMIX_LOG_SOURCE` "pmix.log.source" (`pmix_proc_t*`)

ID of source of the log request.

`PMIX_LOG_STDERR` "pmix.log.stderr" (`char*`)

Log string to `stderr`.

`PMIX_LOG_STDOUT` "pmix.log.stdout" (`char*`)

Log string to `stdout`.

`PMIX_LOG_SYSLOG` "pmix.log.syslog" (`char*`)

Log data to syslog. Defaults to `ERROR` priority. Will log to global syslog if available, otherwise to local syslog.

`PMIX_LOG_LOCAL_SYSLOG` "pmix.log.lsys" (`char*`)

Log data to local syslog. Defaults to `ERROR` priority.

`PMIX_LOG_GLOBAL_SYSLOG` "pmix.log.gsys" (`char*`)

Forward data to system “gateway” and log msg to that syslog Defaults to `ERROR` priority.

`PMIX_LOG_SYSLOG_PRI` "pmix.log.syspri" (`int`)

Syslog priority level.

`PMIX_LOG_TIMESTAMP` "pmix.log.tstmp" (`time_t`)

Timestamp for log report.

`PMIX_LOG_GENERATE_TIMESTAMP` "pmix.log.gtstmp" (`bool`)

Generate timestamp for log.

`PMIX_LOG_TAG_OUTPUT` "pmix.log.tag" (`bool`)

Label the output stream with the channel name (e.g., “stdout”).

`PMIX_LOG_TIMESTAMP_OUTPUT` "pmix.log.tsout" (`bool`)

Print timestamp in output string.

`PMIX_LOG_XML_OUTPUT` "pmix.log.xml" (`bool`)

Print the output stream in XML format.

`PMIX_LOG_ONCE` "pmix.log.once" (`bool`)

Only log this once with whichever channel can first support it, taking the channels in priority order.

`PMIX_LOG_MSG` "pmix.log.msg" (`pmix_byte_object_t`)

1 Message blob to be sent somewhere.

2 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)

3 Log via email based on pmix_info_t containing directives.

4 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)

5 Comma-delimited list of email addresses that are to receive the message.

6 **PMIX_LOG_EMAIL_SENDER_ADDR** "pmix.log.emfaddr" (char*)

7 Return email address of sender.

8 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)

9 Subject line for email.

10 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)

11 Message to be included in email.

12 **PMIX_LOG_EMAIL_SERVER** "pmix.log.esrvr" (char*)

13 Hostname (or IP address) of SMTP server.

14 **PMIX_LOG_EMAIL_SRVR_PORT** "pmix.log.esrvrprt" (int32_t)

15 Port the email server is listening to.

16 **PMIX_LOG_GLOBAL_DATASTORE** "pmix.log.gstore" (bool)

17 Store the log data in a global data store (e.g., database).

18 **PMIX_LOG_JOB_RECORD** "pmix.log.jrec" (bool)

19 Log the provided information to the host environment's job record.

CHAPTER 14

Process Sets and Groups

PMIx supports two slightly related, but functionally different concepts known as *process sets* and *process groups*. This chapter defines these two concepts and describes how they are utilized, along with their corresponding APIs.

14.1 Process Sets

A PMIx *Process Set* is a user-provided or host environment assigned label associated with a given set of application processes. Processes can belong to multiple *process sets* at a time. Users may define a PMIx process set at time of application execution. For example, if using the command line parallel launcher "prun", one could specify process sets as follows:

```
$ prun -n 4 --pset ocean myoceanapp : -n 3 --pset ice myiceapp
```

In this example, the processes in the first application will be labeled with a `PMIX_PSET_NAMES` attribute with a value of `ocean` while those in the second application will be labeled with an `ice` value. During the execution, application processes could lookup the process set attribute for any process using `PMIx_Get`. Alternatively, other executing applications could utilize the `PMIx_Query_info` APIs to obtain the number of declared process sets in the system, a list of their names, and other information about them. In other words, the *process set* identifier provides a label by which an application can derive information about a process and its application - it does *not*, however, confer any operational function.

Host environments can create or delete process sets at any time through the `PMIx_server_define_process_set` and `PMIx_server_delete_process_set` APIs. PMIx servers shall notify all local clients of process set operations via the `PMIX_PROCESS_SET_DEFINE` or `PMIX_PROCESS_SET_DELETE` events.

Process sets differ from *process groups* in several key ways:

- *Process sets* have no implied relationship between their members - i.e., a process in a process set has no concept of a "pset rank" as it would in a process *group*.
- *Process set* identifiers are set by the host environment or by the user at time of application submission for execution - there are no PMIx APIs provided by which an application can define a process set or change a process *set* membership. In contrast, PMIx process *groups* can only be defined dynamically by the application.

- Process *sets* are immutable - members cannot be added or removed once the set has been defined. In contrast, PMIx process *groups* can dynamically change their membership using the appropriate APIs.
- Process *groups* can be used in calls to PMIx operations. Members of process *groups* that are involved in an operation are translated by their PMIx server into their *native* identifier prior to the operation being passed to the host environment. For example, an application can define a process group to consist of ranks 0 and 1 from the host-assigned namespace of 210456, identified by the group id of *foo*. If the application subsequently calls the **PMIx_Fence** API with a process identifier of {*foo*, **PMIX_RANK_WILDCARD**}, the PMIx server will replace that identifier with an array consisting of {210456, 0} and {210456, 1} - the host-assigned identifiers of the participating processes - prior to processing the request.
- Process *groups* can request that the host environment assign a unique **size_t** Process Group Context Identifier (PGCID) to the group at time of group construction. An Message Passing Interface (MPI) library may, for example, use the PGCID as the MPI communicator identifier for the group.

The two concepts do, however, overlap in that they both involve collections of processes. Users desiring to create a process group based on a process set could, for example, obtain the membership array of the process set and use that as input to **PMIx_Group_construct**, perhaps including the process set name as the group identifier for clarity. Note that no linkage between the set and group of the same name is implied nor maintained - e.g., changes in process group membership can not be reflected in the process set using the same identifier.

Advice to PMIx server hosts

The host environment is responsible for ensuring:

- consistent knowledge of process set membership across all involved PMIx servers; and
- that process set names do not conflict with system-assigned namespaces within the scope of the set.

14.1.1 Process Set Constants

PMIx v4.0

The PMIx server is required to send a notification to all local clients upon creation or deletion of process sets. Client processes wishing to receive such notifications must register for the corresponding event:

PMIX_PROCESS_SET_DEFINE -55 The host environment has defined a new process set - the event will include the process set name (**PMIX_PSET_NAME**) and the membership (**PMIX_PSET_MEMBERS**).

PMIX_PROCESS_SET_DELETE -56 The host environment has deleted a process set - the event will include the process set name (**PMIX_PSET_NAME**).

1 14.1.2 Process Set Attributes

2 Several attributes are provided for querying the system regarding process sets using the
3 [PMIx_Query_info](#) APIs.

4 **PMIX_QUERY_NUM_PSETS** "pmix.qry.psetnum" (size_t)

5 Return the number of process sets defined in the specified range (defaults to
6 [PMIX_RANGE_SESSION](#)).

7 **PMIX_QUERY_PSET_NAMES** "pmix.qry.psets" (pmix_data_array_t*)

8 Return a [pmix_data_array_t](#) containing an array of strings of the process set names
9 defined in the specified range (defaults to [PMIX_RANGE_SESSION](#)).

10 **PMIX_QUERY_PSET_MEMBERSHIP** "pmix.qry.pmems" (pmix_data_array_t*)

11 Return an array of [pmix_proc_t](#) containing the members of the specified process set.

12 The [PMIX_PROCESS_SET_DEFINE](#) event shall include the name of the newly defined process
13 set and its members: **PMIX_PSET_NAME** "pmix.pset.nm" (char*)

14 The name of the newly defined process set.

15 **PMIX_PSET_MEMBERS** "pmix.pset.mems" (pmix_data_array_t*)

16 An array of [pmix_proc_t](#) containing the members of the newly defined process set.

17 In addition, a process can request (via [PMIx_Get](#)) the process sets to which a given process
18 (including itself) belongs:

19 **PMIX_PSET_NAMES** "pmix.pset.nms" (pmix_data_array_t*)

20 Returns an array of **char*** string names of the process sets in which the given process is a
21 member.

22 14.2 Process Groups

23 *PMIx Groups* are defined as a collection of processes desiring a common, unique identifier for
24 operational purposes such as passing events or participating in [PMIx fence](#) operations. As with
25 processes that assemble via [PMIx_Connect](#), each member of the group is provided with both the
26 job-level information of any other namespace represented in the group, and the contact information
27 for all group members.

28 However, members of *PMIx Groups* are *loosely coupled* as opposed to *tightly connected* when
29 constructed via [PMIx_Connect](#). Thus, *groups* differ from [PMIx_Connect](#) assemblages in
30 several key areas, as detailed in the following sections.

31 14.2.1 Relation to the host environment

32 Calls to *PMIx Group* APIs are first processed within the local *PMIx* server. When constructed, the
33 server creates a tracker that associates the specified processes with the user-provided group
34 identifier, and assigns a new *group rank* based on their relative position in the array of processes
35 provided in the call to [PMIx_Group_construct](#). Members of the group can subsequently

1 utilize the group identifier in PMIx function calls to address the group’s members, using either
2 **PMIx_RANK_WILDCARD** to refer to all of them or the group-level rank of specific members. The
3 PMIx server will translate the specified processes into their RM-assigned identifiers prior to
4 passing the request up to its host. Thus, the host environment has no visibility into the group’s
5 existence or membership.

6 In contrast, calls to **PMIx_Connect** are relayed to the host environment. This means that the host
7 RM should treat the failure of any process in the specified assemblage as a reportable event and
8 take appropriate action. However, the environment is not required to define a new identifier for the
9 connected assemblage or any of its member processes, nor does it define a new rank for each
10 process within that assemblage. In addition, the PMIx server does not provide any tracking support
11 for the assemblage. Thus, the caller is responsible for addressing members of the connected
12 assemblage using their RM-provided identifiers.

Advice to users

13 User-provided group identifiers must be distinct from both other group identifiers within the system
14 and namespaces provided by the RM so as to avoid collisions between group identifiers and
15 RM-assigned namespaces. This can usually be accomplished through the use of an
16 application-specific prefix – e.g., “myapp-foo”

17 14.2.2 Construction procedure

18 **PMIx_Connect** calls require that every process call the API before completing – i.e., it is
19 modeled upon the bulk synchronous traditional MPI connect/accept methodology. Thus, a given
20 application thread can only be involved in one connect/accept operation at a time, and is blocked in
21 that operation until all specified processes participate. In addition, there is no provision for
22 replacing processes in the assemblage due to failure to participate, nor a mechanism by which a
23 process might decline participation.

24 In contrast, PMIx Groups are designed to be more flexible in their construction procedure by
25 relaxing these constraints. While a standard blocking form of constructing groups is provided, the
26 event notification system is utilized to provide a designated *group leader* with the ability to replace
27 participants that fail to participate within a given timeout period. This provides a mechanism by
28 which the application can, if desired, replace members on-the-fly or allow the group to proceed
29 with partial membership. In such cases, the final group membership is returned to all participants
30 upon completion of the operation.

31 Additionally, PMIx supports dynamic definition of group membership based on an invite/join
32 model. A process can asynchronously initiate construction of a group of any processes via the
33 **PMIx_Group_invite** function call. Invitations are delivered via a PMIx event (using the
34 **PMIx_GROUP_INVITED** event) to the invited processes which can then either accept or decline
35 the invitation using the **PMIx_Group_join** API. The initiating process tracks responses by
36 registering for the events generated by the call to **PMIx_Group_join**, timeouts, or process

1 terminations, optionally replacing processes that decline the invitation, fail to respond in time, or
2 terminate without responding. Upon completion of the operation, the final list of participants is
3 communicated to each member of the new group.

4 14.2.3 Destruct procedure

5 Members of a PMIx Group may depart the group at any time via the `PMIx_Group_leave` API.
6 Other members are notified of the departure via the `PMIX_GROUP_LEFT` event to distinguish such
7 events from those reporting process termination. This leaves the remaining members free to
8 continue group operations. The `PMIx_Group_destruct` operation offers a collective method
9 akin to `PMIx_Disconnect` for deconstructing the entire group.

10 In contrast, processes that assemble via `PMIx_Connect` must all depart the assemblage together –
11 i.e., no member can depart the assemblage while leaving the remaining members in it. Even the
12 non-blocking form of `PMIx_Disconnect` retains this requirement in that members remain a part
13 of the assemblage until all members have called `PMIx_Disconnect_nb`

14 Note that applications supporting dynamic group behaviors such as asynchronous departure take
15 responsibility for ensuring global consistency in the group definition prior to executing group
16 collective operations - i.e., it is the application's responsibility to either ensure that knowledge of
17 the current group membership is globally consistent across the participants, or to register for
18 appropriate events to deal with the lack of consistency during the operation.

▼ Advice to users ▼

19 The reliance on PMIx events in the PMIx Group concept dictates that processes utilizing these APIs
20 must register for the corresponding events. Failure to do so will likely lead to operational failures.
21 Users are recommended to utilize the `PMIX_TIMEOUT` directive (or retain an internal timer) on
22 calls to PMIx Group APIs (especially the blocking form of those functions) as processes that have
23 not registered for required events will never respond.

24 14.2.4 Process Group Events

25 *PMIx v4.0*

26 Asynchronous process group operations rely heavily on PMIx events. The following events have
27 been defined for that purpose.

28 **PMIX_GROUP_INVITED -159** The process has been invited to join a PMIx Group - the
29 identifier of the group and the ID's of other invited (or already joined) members will be
30 included in the notification.

31 **PMIX_GROUP_LEFT -160** A process has asynchronously left a PMIx Group - the process
32 identifier of the departing process will in included in the notification.

33 **PMIX_GROUP_MEMBER_FAILED -170** A member of a PMIx Group has abnormally
34 terminated (i.e., without formally leaving the group prior to termination) - the process
35 identifier of the failed process will be included in the notification.

1 **PMIX_GROUP_INVITE_ACCEPTED** -161 A process has accepted an invitation to join a
 2 PMIx Group - the identifier of the group being joined will be included in the notification.
 3 **PMIX_GROUP_INVITE_DECLINED** -162 A process has declined an invitation to join a
 4 PMIx Group - the identifier of the declined group will be included in the notification.
 5 **PMIX_GROUP_INVITE_FAILED** -163 An invited process failed or terminated prior to
 6 responding to the invitation - the identifier of the failed process will be included in the
 7 notification.
 8 **PMIX_GROUP_MEMBERSHIP_UPDATE** -164 The membership of a PMIx group has
 9 changed - the identifiers of the revised membership will be included in the notification.
 10 **PMIX_GROUP_CONSTRUCT_ABORT** -165 Any participant in a PMIx group construct
 11 operation that returns **PMIX_GROUP_CONSTRUCT_ABORT** from the *leader failed* event
 12 handler will cause all participants to receive an event notifying them of that status. Similarly,
 13 the leader may elect to abort the procedure by either returning this error code from the handler
 14 assigned to the **PMIX_GROUP_INVITE_ACCEPTED** or
 15 **PMIX_GROUP_INVITE_DECLINED** codes, or by generating an event for the abort code.
 16 Abort events will be sent to all invited or existing members of the group.
 17 **PMIX_GROUP_CONSTRUCT_COMPLETE** -166 The group construct operation has
 18 completed - the final membership will be included in the notification.
 19 **PMIX_GROUP_LEADER_FAILED** -168 The current *leader* of a group including this
 20 process has abnormally terminated - the group identifier will be included in the notification.
 21 **PMIX_GROUP_LEADER_SELECTED** -167 A new *leader* of a group including this process
 22 has been selected - the identifier of the new leader will be included in the notification.
 23 **PMIX_GROUP_CONTEXT_ID_ASSIGNED** -169 A new PGCID has been assigned by the
 24 host environment to a group that includes this process - the group identifier will be included in
 25 the notification.

26 14.2.5 Process Group Attributes

27 *PMIx v4.0*

Attributes for querying the system regarding process groups include:

28 **PMIX_QUERY_NUM_GROUPS** "pmix.qry.pgrpnum" (**size_t**)
 29 Return the number of process groups defined in the specified range (defaults to session).
 30 OPTIONAL QUALIFIERS: **PMIX_RANGE**.
 31 **PMIX_QUERY_GROUP_NAMES** "pmix.qry.pgrp" (**pmix_data_array_t***)
 32 Return a **pmix_data_array_t** containing an array of string names of the process groups
 33 defined in the specified range (defaults to session). OPTIONAL QUALIFIERS:
 34 **PMIX_RANGE**.
 35 **PMIX_QUERY_GROUP_MEMBERSHIP**
 36 "pmix.qry.pgrpmems" (**pmix_data_array_t***)
 37 Return a **pmix_data_array_t** of **pmix_proc_t** containing the members of the
 38 specified process group. REQUIRED QUALIFIERS: **PMIX_GROUP_ID**.

39 The following attributes are used as directives in PMIx Group operations:

40 **PMIX_GROUP_ID** "pmix.grp.id" (**char***)

1 User-provided group identifier - as the group identifier may be used in PMIx operations, the
2 user is required to ensure that the provided ID is unique within the scope of the host
3 environment (e.g., by including some user-specific or application-specific prefix or suffix to
4 the string).

5 **PMIX_GROUP_LEADER** "pmix.grp.ldr" (bool)

6 This process is the leader of the group.

7 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (bool)

8 Participation is optional - do not return an error if any of the specified processes terminate
9 without having joined. The default is **false**.

10 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)

11 Notify remaining members when another member terminates without first leaving the group.

12 **PMIX_GROUP_FT_COLLECTIVE** "pmix.grp.ftcoll" (bool)

13 Adjust internal tracking on-the-fly for terminated processes during a PMIx group collective
14 operation.

15 **PMIX_GROUP_MEMBERSHIP** "pmix.grp.mbrs" (pmix_data_array_t*)

16 Array **pmix_proc_t** identifiers identifying the members of the specified group.

17 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)

18 Requests that the RM assign a new context identifier to the newly created group. The
19 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
20 specified in the request. Thus, the value serves as a means of identifying the group within
21 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

22 **PMIX_GROUP_LOCAL_ONLY** "pmix.grp.lcl" (bool)

23 Group operation only involves local processes. PMIx implementations are *required* to
24 automatically scan an array of group members for local vs remote processes - if only local
25 processes are detected, the implementation need not execute a global collective for the
26 operation unless a context ID has been requested from the host environment. This can result
27 in significant time savings. This attribute can be used to optimize the operation by indicating
28 whether or not only local processes are represented, thus allowing the implementation to
29 bypass the scan.

30 The following attributes are used to return information at the conclusion of a PMIx Group
31 operation and/or in event notifications:

32 **PMIX_GROUP_CONTEXT_ID** "pmix.grp.ctxid" (size_t)

33 Context identifier assigned to the group by the host RM.

34 **PMIX_GROUP_ENDPT_DATA** "pmix.grp.endpt" (pmix_byte_object_t)

35 Data collected during group construction to ensure communication between group members
36 is supported upon completion of the operation.

37 In addition, a process can request (via **PMIx_Get**) the process groups to which a given process
38 (including itself) belongs:

39 **PMIX_GROUP_NAMES** "pmix.pgrp.nm" (pmix_data_array_t*)

1 Returns an array of **char*** string names of the process groups in which the given process is
2 a member.

3 14.2.6 PMIx_Group_construct

4 Summary

5 Construct a PMIx process group.

6 Format

C

PMIx v4.0

```
7 pmix_status_t  
8 PMIx_Group_construct(const char grp[],  
9                     const pmix_proc_t procs[], size_t nprocs,  
10                    const pmix_info_t directives[],  
11                    size_t ndirs,  
12                    pmix_info_t **results,  
13                    size_t *nresults);
```

C

- 14 **IN grp**
15 **NULL**-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the group
16 identifier (string)
- 17 **IN procs**
18 Array of **pmix_proc_t** structures containing the PMIx identifiers of the member processes
19 (array of handles)
- 20 **IN nprocs**
21 Number of elements in the *procs* array (**size_t**)
- 22 **IN directives**
23 Array of **pmix_info_t** structures (array of handles)
- 24 **IN ndirs**
25 Number of elements in the *directives* array (**size_t**)
- 26 **INOUT results**
27 Pointer to a location where the array of **pmix_info_t** describing the results of the
28 operation is to be returned (pointer to handle)
- 29 **INOUT nresults**
30 Pointer to a **size_t** location where the number of elements in *results* is to be returned
31 (memory reference)
- 32 Returns **PMIX_SUCCESS** or a negative value indicating the error.

Required Attributes

The following attributes are *required* to be supported by all PMIx libraries that support this operation:

PMIX_GROUP_LEADER "pmix.grp.ldr" (bool)

This process is the leader of the group.

PMIX_GROUP_OPTIONAL "pmix.grp.opt" (bool)

Participation is optional - do not return an error if any of the specified processes terminate without having joined. The default is **false**.

PMIX_GROUP_LOCAL_ONLY "pmix.grp.lcl" (bool)

Group operation only involves local processes. PMIx implementations are *required* to automatically scan an array of group members for local vs remote processes - if only local processes are detected, the implementation need not execute a global collective for the operation unless a context ID has been requested from the host environment. This can result in significant time savings. This attribute can be used to optimize the operation by indicating whether or not only local processes are represented, thus allowing the implementation to bypass the scan.

PMIX_GROUP_FT_COLLECTIVE "pmix.grp.ftcoll" (bool)

Adjust internal tracking on-the-fly for terminated processes during a PMIx group collective operation.

Host environments that support this operation are *required* to support the following attributes:

PMIX_GROUP_ASSIGN_CONTEXT_ID "pmix.grp.actxid" (bool)

Requests that the RM assign a new context identifier to the newly created group. The identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range specified in the request. Thus, the value serves as a means of identifying the group within that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

PMIX_GROUP_NOTIFY_TERMINATION "pmix.grp.notterm" (bool)

Notify remaining members when another member terminates without first leaving the group.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Construct a new group composed of the specified processes and identified with the provided group identifier. The group identifier is a user-defined, **NULL**-terminated character array of length less than or equal to **PMIX_MAX_NSLEN**. Only characters accepted by standard string comparison functions (e.g., *strncmp*) are supported. Processes may engage in multiple simultaneous group construct operations so long as each is provided with a unique group ID. The *directives* array can be used to pass user-level directives regarding timeout constraints and other options available from the PMIx server.

If the **PMIX_GROUP_NOTIFY_TERMINATION** attribute is provided and has a value of **true**, then either the construct leader (if **PMIX_GROUP_LEADER** is provided) or all participants who register for the **PMIX_GROUP_MEMBER_FAILED** event will receive events whenever a process fails or terminates prior to calling **PMIx_Group_construct** – i.e. if a *group leader* is declared, *only* that process will receive the event. In the absence of a declared leader, *all* specified group members will receive the event.

The event will contain the identifier of the process that failed to join plus any other information that the host RM provided. This provides an opportunity for the leader or the collective members to react to the event – e.g., to decide to proceed with a smaller group or to abort the operation. The decision is communicated to the PMIx library in the results array at the end of the event handler. This allows PMIx to properly adjust accounting for procedure completion. When construct is complete, the participating PMIx servers will be alerted to any change in participants and each group member will receive an updated group membership (marked with the **PMIX_GROUP_MEMBERSHIP** attribute) as part of the *results* array returned by this API.

Failure of the declared leader at any time will cause a **PMIX_GROUP_LEADER_FAILED** event to be delivered to all participants so they can optionally declare a new leader. A new leader is identified by providing the **PMIX_GROUP_LEADER** attribute in the results array in the return of the event handler. Only one process is allowed to return that attribute, thereby declaring itself as the new leader. Results of the leader selection will be communicated to all participants via a **PMIX_GROUP_LEADER_SELECTED** event identifying the new leader. If no leader was selected, then the **pmix_info_t** provided to that event handler will include that information so the participants can take appropriate action.

Any participant that returns **PMIX_GROUP_CONSTRUCT_ABORT** from either the **PMIX_GROUP_MEMBER_FAILED** or the **PMIX_GROUP_LEADER_FAILED** event handler will cause the construct process to abort, returning from the call with a **PMIX_GROUP_CONSTRUCT_ABORT** status.

If the **PMIX_GROUP_NOTIFY_TERMINATION** attribute is not provided or has a value of **false**, then the **PMIx_Group_construct** operation will simply return an error whenever a proposed group member fails or terminates prior to calling **PMIx_Group_construct**.

Providing the **PMIX_GROUP_OPTIONAL** attribute with a value of **true** directs the PMIx library to consider participation by any specified group member as non-required - thus, the operation will return **PMIX_SUCCESS** if all members participate, or **PMIX_ERR_PARTIAL_SUCCESS** if some

1 members fail to participate. The *results* array will contain the final group membership in the latter
2 case. Note that this use-case can cause the operation to hang if the `PMIX_TIMEOUT` attribute is
3 not specified and one or more group members fail to call `PMIx_Group_construct` while
4 continuing to execute. Also, note that no leader or member failed events will be generated during
5 the operation.

6 Processes in a group under construction are not allowed to leave the group until group construction
7 is complete. Upon completion of the construct procedure, each group member will have access to
8 the job-level information of all namespaces represented in the group plus any information posted
9 via `PMIx_Put` (subject to the usual scoping directives) for every group member.

Advice to PMIx library implementers

10 At the conclusion of the construct operation, the PMIx library is *required* to ensure that job-related
11 information from each participating namespace plus any information posted by group members via
12 `PMIx_Put` (subject to scoping directives) is available to each member via calls to `PMIx_Get`.

Advice to PMIx server hosts

13 The collective nature of this API generally results in use of a fence-like operation by the backend
14 host environment. Host environments that utilize the array of process participants as a *signature* for
15 such operations may experience potential conflicts should both a `PMIx_Group_construct` and
16 a `PMIx_Fence` operation involving the same participants be simultaneously executed. As PMIx
17 allows for such use-cases, it is therefore the responsibility of the host environment to resolve any
18 potential conflicts.

19 14.2.7 `PMIx_Group_construct_nb`

20 Summary

21 Non-blocking form of `PMIx_Group_construct`.

Format

C

```
pmix_status_t
PMIx_Group_construct_nb(const char grp[],
                       const pmix_proc_t procs[], size_t nprocs,
                       const pmix_info_t directives[],
                       size_t ndirs,
                       pmix_info_cbfunc_t cbfunc, void *cbdata);
```

C

- IN grp**
NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the group identifier (string)
- IN procs**
Array of **pmix_proc_t** structures containing the PMIx identifiers of the member processes (array of handles)
- IN nprocs**
Number of elements in the *procs* array (**size_t**)
- IN directives**
Array of **pmix_info_t** structures (array of handles)
- IN ndirs**
Number of elements in the *directives* array (**size_t**)
- IN cbfunc**
Callback function **pmix_info_cbfunc_t** (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

PMIX_OPERATION_SUCCEEDED, indicating that the request was immediately processed successfully - the *cbfunc* will *not* be called.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

If executed, the status returned in the provided callback function will be one of the following constants:

- **PMIX_SUCCESS** The operation succeeded and all specified members participated.
- **PMIX_ERR_PARTIAL_SUCCESS** The operation succeeded but not all specified members participated - the final group membership is included in the callback function.

- 1 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM
2 does not.
- 3 • a non-zero PMIx error constant indicating a reason for the request's failure.

Required Attributes

4 PMIx libraries that choose not to support this operation *must* return
5 **PMIX_ERR_NOT_SUPPORTED** when the function is called.

6 The following attributes are *required* to be supported by all PMIx libraries that support this
7 operation:

8 **PMIX_GROUP_LEADER** "pmix.grp.ldr" (bool)

9 This process is the leader of the group.

10 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (bool)

11 Participation is optional - do not return an error if any of the specified processes terminate
12 without having joined. The default is **false**.

13 **PMIX_GROUP_LOCAL_ONLY** "pmix.grp.lcl" (bool)

14 Group operation only involves local processes. PMIx implementations are *required* to
15 automatically scan an array of group members for local vs remote processes - if only local
16 processes are detected, the implementation need not execute a global collective for the
17 operation unless a context ID has been requested from the host environment. This can result
18 in significant time savings. This attribute can be used to optimize the operation by indicating
19 whether or not only local processes are represented, thus allowing the implementation to
20 bypass the scan.

21 **PMIX_GROUP_FT_COLLECTIVE** "pmix.grp.ftcoll" (bool)

22 Adjust internal tracking on-the-fly for terminated processes during a PMIx group collective
23 operation.

24 Host environments that support this operation are *required* to provide the following attributes:

25 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)

26 Requests that the RM assign a new context identifier to the newly created group. The
27 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
28 specified in the request. Thus, the value serves as a means of identifying the group within
29 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

30 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)

31 Notify remaining members when another member terminates without first leaving the group.
32

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Non-blocking version of the **PMIx_Group_construct** operation. The callback function will be called once all group members have called either **PMIx_Group_construct** or **PMIx_Group_construct_nb**.

14.2.8 PMIx_Group_destruct

Summary

Destruct a PMIx process group.

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Group_destruct(const char grp[],
                   const pmix_info_t directives[],
                   size_t ndirs);
```

C

IN grp

NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the identifier of the group to be destructed (string)

IN directives

Array of **pmix_info_t** structures (array of handles)

IN ndirs

Number of elements in the *directives* array (**size_t**)

Returns **PMIX_SUCCESS** or a negative value indicating the error.

Required Attributes

For implementations and host environments that support the operation, there are no identified required attributes for this API.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Destruct a group identified by the provided group identifier. Processes may engage in multiple simultaneous group destruct operations so long as each involves a unique group ID. The *directives* array can be used to pass user-level directives regarding timeout constraints and other options available from the PMIx server.

The destruct API will return an error if any group process fails or terminates prior to calling **PMIx_Group_destruct** or its non-blocking version unless the **PMIX_GROUP_NOTIFY_TERMINATION** attribute was provided (with a value of **false**) at time of group construction. If notification was requested, then the **PMIX_GROUP_MEMBER_FAILED** event will be delivered for each process that fails to call destruct and the destruct tracker updated to account for the lack of participation. The **PMIx_Group_destruct** operation will subsequently return **PMIX_SUCCESS** when the remaining processes have all called destruct – i.e., the event will serve in place of return of an error.

Advice to PMIx server hosts

The collective nature of this API generally results in use of a fence-like operation by the backend host environment. Host environments that utilize the array of process participants as a *signature* for such operations may experience potential conflicts should both a **PMIx_Group_destruct** and a **PMIx_Fence** operation involving the same participants be simultaneously executed. As PMIx allows for such use-cases, it is therefore the responsibility of the host environment to resolve any potential conflicts.

14.2.9 PMIx_Group_destruct_nb

Summary

Non-blocking form of **PMIx_Group_destruct**.

Format

C

```
pmix_status_t
PMIx_Group_destruct_nb(const char grp[],
                      const pmix_info_t directives[],
                      size_t ndirs,
                      pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

- IN grp**
NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the identifier of the group to be destructed (string)
- IN directives**
Array of **pmix_info_t** structures (array of handles)
- IN ndirs**
Number of elements in the *directives* array (**size_t**)
- IN cbfunc**
Callback function **pmix_op_cbfunc_t** (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

PMIX_OPERATION_SUCCEEDED, indicating that the request was immediately processed successfully - the *cbfunc* will *not* be called.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

If executed, the status returned in the provided callback function will be one of the following constants:

- **PMIX_SUCCESS** The operation was successfully completed.
- **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM does not.
- a non-zero PMIx error constant indicating a reason for the request's failure.

Required Attributes

PMIx libraries that choose not to support this operation *must* return **PMIX_ERR_NOT_SUPPORTED** when the function is called. For implementations and host environments that support the operation, there are no identified required attributes for this API.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Non-blocking version of the **PMIx_Group_destruct** operation. The callback function will be called once all members of the group have executed either **PMIx_Group_destruct** or **PMIx_Group_destruct_nb**.

14.2.10 PMIx_Group_invite

Summary

Asynchronously construct a PMIx process group.

Format

PMIx v4.0

C

pmix_status_t

```
PMIx_Group_invite(const char grp[],  
                 const pmix_proc_t procs[], size_t nprocs,  
                 const pmix_info_t directives[], size_t ndirs,  
                 pmix_info_t **results, size_t *nresult);
```

C

IN grp

NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the group identifier (string)

IN procs

Array of **pmix_proc_t** structures containing the PMIx identifiers of the processes to be invited (array of handles)

IN nprocs

Number of elements in the *procs* array (**size_t**)

IN directives

Array of **pmix_info_t** structures (array of handles)

IN ndirs

Number of elements in the *directives* array (**size_t**)

1 **INOUT results**

2 Pointer to a location where the array of `pmix_info_t` describing the results of the
3 operation is to be returned (pointer to handle)

4 **INOUT nresults**

5 Pointer to a `size_t` location where the number of elements in *results* is to be returned
6 (memory reference)

7 Returns `PMIX_SUCCESS` or a negative value indicating the error.

8 **Required Attributes**

9 The following attributes are *required* to be supported by all PMIx libraries that support this
10 operation:

11 **PMIX_GROUP_OPTIONAL** "`pmix.grp.opt`" (`bool`)

12 Participation is optional - do not return an error if any of the specified processes terminate
13 without having joined. The default is `false`.

14 **PMIX_GROUP_FT_COLLECTIVE** "`pmix.grp.ftcoll`" (`bool`)

15 Adjust internal tracking on-the-fly for terminated processes during a PMIx group collective
16 operation.

17 Host environments that support this operation are *required* to provide the following attributes:

18 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "`pmix.grp.actxid`" (`bool`)

19 Requests that the RM assign a new context identifier to the newly created group. The
20 identifier is an unsigned, `size_t` value that the RM guarantees to be unique across the range
21 specified in the request. Thus, the value serves as a means of identifying the group within
22 that range. If no range is specified, then the request defaults to `PMIX_RANGE_SESSION`.

23 **PMIX_GROUP_NOTIFY_TERMINATION** "`pmix.grp.notterm`" (`bool`)

24 Notify remaining members when another member terminates without first leaving the group.

25 **Optional Attributes**

26 The following attributes are optional for host environments that support this operation:

27 **PMIX_TIMEOUT** "`pmix.timeout`" (`int`)

28 Time in seconds before the specified operation should time out (zero indicating infinite) and
29 return the `PMIX_ERR_TIMEOUT` error. Care should be taken to avoid race conditions
caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Explicitly invite the specified processes to join a group. The process making the `PMIx_Group_invite` call is automatically declared to be the *group leader*. Each invited process will be notified of the invitation via the `PMIX_GROUP_INVITED` event - the processes being invited must therefore register for the `PMIX_GROUP_INVITED` event in order to be notified of the invitation. Note that the PMIx event notification system caches events - thus, no ordering of invite versus event registration is required.

The invitation event will include the identity of the inviting process plus the name of the group. When ready to respond, each invited process provides a response using either the blocking or non-blocking form of `PMIx_Group_join`. This will notify the inviting process that the invitation was either accepted (via the `PMIX_GROUP_INVITE_ACCEPTED` event) or declined (via the `PMIX_GROUP_INVITE_DECLINED` event). The `PMIX_GROUP_INVITE_ACCEPTED` event is captured by the PMIx client library of the inviting process - i.e., the application itself does not need to register for this event. The library will track the number of accepting processes and alert the inviting process (by returning from the blocking form of `PMIx_Group_invite` or calling the callback function of the non-blocking form) when group construction completes.

The inviting process should, however, register for the `PMIX_GROUP_INVITE_DECLINED` if the application allows invited processes to decline the invitation. This provides an opportunity for the application to either invite a replacement, declare “abort”, or choose to remove the declining process from the final group. The inviting process should also register to receive `PMIX_GROUP_INVITE_FAILED` events whenever a process fails or terminates prior to responding to the invitation. Actions taken by the inviting process in response to these events must be communicated at the end of the event handler by returning the corresponding result so that the PMIx library can adjust accordingly.

Upon completion of the operation, all members of the new group will receive access to the job-level information of each other’s namespaces plus any information posted via `PMIx_Put` by the other members.

The inviting process is automatically considered the leader of the asynchronous group construction procedure and will receive all failure or termination events for invited members prior to completion. The inviting process is required to provide a `PMIX_GROUP_CONSTRUCT_COMPLETE` event once the group has been fully assembled - this event is used by the PMIx library as a trigger to release participants from their call to `PMIx_Group_join` and provides information (e.g., the final group membership) to be returned in the *results* array.

Failure of the inviting process at any time will cause a `PMIX_GROUP_LEADER_FAILED` event to be delivered to all participants so they can optionally declare a new leader. A new leader is identified by providing the `PMIX_GROUP_LEADER` attribute in the results array in the return of the event handler. Only one process is allowed to return that attribute, declaring itself as the new leader. Results of the leader selection will be communicated to all participants via a `PMIX_GROUP_LEADER_SELECTED` event identifying the new leader. If no leader was selected, then the status code provided in the event handler will provide an error value so the participants can take appropriate action.

Advice to users

Applications are not allowed to use the group in any operations until group construction is complete. This is required in order to ensure consistent knowledge of group membership across all participants.

14.2.11 PMIx_Group_invite_nb

Summary

Non-blocking form of `PMIx_Group_invite`.

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Group_invite_nb(const char grp[],
                    const pmix_proc_t procs[], size_t nprocs,
                    const pmix_info_t directives[], size_t ndirs,
                    pmix_info_cbfunc_t cbfunc, void *cbdata);
```

C

IN grp

NULL-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the group identifier (string)

IN procs

Array of `pmix_proc_t` structures containing the PMIx identifiers of the processes to be invited (array of handles)

IN nprocs

Number of elements in the *procs* array (`size_t`)

IN directives

Array of `pmix_info_t` structures (array of handles)

IN ndirs

Number of elements in the *directives* array (`size_t`)

IN cbfunc

Callback function `pmix_info_cbfunc_t` (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when `PMIX_SUCCESS` is returned.

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

1 **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed
2 successfully - the *cbfunc* will *not* be called.

3 If none of the above return codes are appropriate, then an implementation must return either a
4 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

5 If executed, the status returned in the provided callback function will be one of the following
6 constants:

- 7 • **PMIX_SUCCESS** The operation succeeded and all specified members participated.
- 8 • **PMIX_ERR_PARTIAL_SUCCESS** The operation succeeded but not all specified members
9 participated - the final group membership is included in the callback function.
- 10 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM
11 does not.
- 12 • a non-zero PMIx error constant indicating a reason for the request's failure.

▼----- Required Attributes -----▼

13 The following attributes are *required* to be supported by all PMIx libraries that support this
14 operation:

15 **PMIX_GROUP_OPTIONAL** "**pmix.grp.opt**" (**bool**)

16 Participation is optional - do not return an error if any of the specified processes terminate
17 without having joined. The default is **false**.

18 **PMIX_GROUP_FT_COLLECTIVE** "**pmix.grp.ftcoll**" (**bool**)

19 Adjust internal tracking on-the-fly for terminated processes during a PMIx group collective
20 operation.

21 Host environments that support this operation are *required* to provide the following attributes:

22 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "**pmix.grp.actxid**" (**bool**)

23 Requests that the RM assign a new context identifier to the newly created group. The
24 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
25 specified in the request. Thus, the value serves as a means of identifying the group within
26 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

27 **PMIX_GROUP_NOTIFY_TERMINATION** "**pmix.grp.notterm**" (**bool**)

28 Notify remaining members when another member terminates without first leaving the group.
29

▲-----▲

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Non-blocking version of the **PMIx_Group_invite** operation. The callback function will be called once all invited members of the group (or their substitutes) have executed either **PMIx_Group_join** or **PMIx_Group_join_nb**.

14.2.12 PMIx_Group_join

Summary

Accept an invitation to join a PMIx process group.

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Group_join(const char grp[],
                const pmix_proc_t *leader,
                pmix_group_opt_t opt,
                const pmix_info_t directives[], size_t ndirs,
                pmix_info_t **results, size_t *nresult);
```

C

IN grp

NULL-terminated character array of maximum size **PMIX_MAX_NSLEN** containing the group identifier (string)

IN leader

Process that generated the invitation (handle)

IN opt

Accept or decline flag (**pmix_group_opt_t**)

IN directives

Array of **pmix_info_t** structures (array of handles)

IN ndirs

Number of elements in the *directives* array (**size_t**)

INOUT results

Pointer to a location where the array of **pmix_info_t** describing the results of the operation is to be returned (pointer to handle)

1 **INOUT `nresults`**

2 Pointer to a `size_t` location where the number of elements in `results` is to be returned
3 (memory reference)

4 Returns `PMIX_SUCCESS` or a negative value indicating the error.

▼----- Required Attributes -----▼

5 There are no identified required attributes for implementers.



▼----- Optional Attributes -----▼

6 The following attributes are optional for host environments that support this operation:

7 `PMIX_TIMEOUT` "`pmix.timeout`" (`int`)

8 Time in seconds before the specified operation should time out (zero indicating infinite) and
9 return the `PMIX_ERR_TIMEOUT` error. Care should be taken to avoid race conditions
10 caused by multiple layers (client, server, and host) simultaneously timing the operation.



11 **Description**

12 Respond to an invitation to join a group that is being asynchronously constructed. The process must
13 have registered for the `PMIX_GROUP_INVITED` event in order to be notified of the invitation.
14 When called, the event information will include the `pmix_proc_t` identifier of the process that
15 generated the invitation along with the identifier of the group being constructed. When ready to
16 respond, the process provides a response using either form of `PMIx_Group_join`.

▼----- Advice to users -----▼

17 Since the process is alerted to the invitation in a PMIx event handler, the process *must not* use the
18 blocking form of this call unless it first “thread shifts” out of the handler and into its own thread
19 context. Likewise, while it is safe to call the non-blocking form of the API from the event handler,
20 the process *must not* block in the handler while waiting for the callback function to be called.

1 Calling this function causes the inviting process (aka the *group leader*) to be notified that the
2 process has either accepted or declined the request. The blocking form of the API will return once
3 the group has been completely constructed or the group's construction has failed (as described
4 below) – likewise, the callback function of the non-blocking form will be executed upon the same
5 conditions.

6 Failure of the leader during the call to `PMIx_Group_join` will cause a
7 `PMIX_GROUP_LEADER_FAILED` event to be delivered to all invited participants so they can
8 optionally declare a new leader. A new leader is identified by providing the
9 `PMIX_GROUP_LEADER` attribute in the results array in the return of the event handler. Only one
10 process is allowed to return that attribute, declaring itself as the new leader. Results of the leader
11 selection will be communicated to all participants via a `PMIX_GROUP_LEADER_SELECTED`
12 event identifying the new leader. If no leader was selected, then the status code provided in the
13 event handler will provide an error value so the participants can take appropriate action.

14 Any participant that returns `PMIX_GROUP_CONSTRUCT_ABORT` from the leader failed event
15 handler will cause all participants to receive an event notifying them of that status. Similarly, the
16 leader may elect to abort the procedure by either returning `PMIX_GROUP_CONSTRUCT_ABORT`
17 from the handler assigned to the `PMIX_GROUP_INVITE_ACCEPTED` or
18 `PMIX_GROUP_INVITE_DECLINED` codes, or by generating an event for the abort code. Abort
19 events will be sent to all invited participants.

20 14.2.13 `PMIx_Group_join_nb`

21 Summary

22 Non-blocking form of `PMIx_Group_join`

23 Format

PMIx v4.0

C

24 `pmix_status_t`

```
25 PMIx_Group_join_nb(const char grp[],  
26                   const pmix_proc_t *leader,  
27                   pmix_group_opt_t opt,  
28                   const pmix_info_t directives[], size_t ndirs,  
29                   pmix_info_cbfunc_t cbfunc, void *cbdata);
```

C

30 **IN** `grp`

31 `NULL`-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the group
32 identifier (string)

33 **IN** `leader`

34 Process that generated the invitation (handle)

1 **IN** **opt**
 2 Accept or decline flag (`pmix_group_opt_t`)
 3 **IN** **directives**
 4 Array of `pmix_info_t` structures (array of handles)
 5 **IN** **ndirs**
 6 Number of elements in the *directives* array (`size_t`)
 7 **IN** **cbfunc**
 8 Callback function `pmix_info_cbfunc_t` (function reference)
 9 **IN** **cbdata**
 10 Data to be passed to the callback function (memory reference)

11 A successful return indicates that the request is being processed and the result will be returned in
 12 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
 13 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

14 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

15 **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed
 16 successfully - the *cbfunc* will *not* be called.

17 If none of the above return codes are appropriate, then an implementation must return either a
 18 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

19 If executed, the status returned in the provided callback function will be one of the following
 20 constants:

- 21 • **PMIX_SUCCESS** The operation succeeded and group membership is in the callback function
 22 parameters.
- 23 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx server supports this operation, the host RM
 24 does not.
- 25 • a non-zero PMIx error constant indicating a reason for the request's failure.

▼----- Required Attributes -----▼

26 There are no identified required attributes for implementers.

▲-----

▼----- Optional Attributes -----▼

27 The following attributes are optional for host environments that support this operation:

28 **PMIX_TIMEOUT** "`pmix.timeout`" (`int`)
 29 Time in seconds before the specified operation should time out (zero indicating infinite) and
 30 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
 31 caused by multiple layers (client, server, and host) simultaneously timing the operation.

▲-----

Description

Non-blocking version of the `PMIx_Group_join` operation. The callback function will be called once all invited members of the group (or their substitutes) have executed either `PMIx_Group_join` or `PMIx_Group_join_nb`.

14.2.13.1 Group accept/decline directives

PMIx v4.0

The `pmix_group_opt_t` type is a `uint8_t` value used with the `PMIx_Group_join` API to indicate *accept* or *decline* of the invitation - these are provided for readability of user code:

`PMIX_GROUP_DECLINE` 0 Decline the invitation.

`PMIX_GROUP_ACCEPT` 1 Accept the invitation.

14.2.14 `PMIx_Group_leave`

Summary

Leave a `PMIx` process group.

Format

PMIx v4.0

```
pmix_status_t
PMIx_Group_leave(const char grp[],
                 const pmix_info_t directives[],
                 size_t ndirs);
```

IN `grp`

`NULL`-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the group identifier (string)

IN `directives`

Array of `pmix_info_t` structures (array of handles)

IN `ndirs`

Number of elements in the *directives* array (`size_t`)

Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

There are no identified required attributes for implementers.

Description

Calls to `PMIx_Group_leave` (or its non-blocking form) will cause a `PMIX_GROUP_LEFT` event to be generated notifying all members of the group of the caller's departure. The function will return (or the non-blocking function will execute the specified callback function) once the event has been locally generated and is not indicative of remote receipt.

Advice to users

The `PMIx_Group_leave` API is intended solely for asynchronous departures of individual processes from a group as it is not a scalable operation – i.e., when a process determines it should no longer be a part of a defined group, but the remainder of the group retains a valid reason to continue in existence. Developers are advised to use `PMIx_Group_destruct` (or its non-blocking form) for all other scenarios as it represents a more scalable operation.

14.2.15 `PMIx_Group_leave_nb`

Summary

Non-blocking form of `PMIx_Group_leave`.

Format

PMIx v4.0

C

```
pmix_status_t
PMIx_Group_leave_nb(const char grp[],
                    const pmix_info_t directives[],
                    size_t ndirs,
                    pmix_op_cbfunc_t cbfunc,
                    void *cbdata);
```

C

- IN** `grp`
NULL-terminated character array of maximum size `PMIX_MAX_NSLEN` containing the group identifier (string)
- IN** `directives`
Array of `pmix_info_t` structures (array of handles)
- IN** `ndirs`
Number of elements in the `directives` array (`size_t`)
- IN** `cbfunc`
Callback function `pmix_op_cbfunc_t` (function reference)
- IN** `cbdata`
Data to be passed to the callback function (memory reference)

1 A successful return indicates that the request is being processed and the result will be returned in
2 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
3 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

4 Returns PMIX_SUCCESS or one of the following error codes when the condition described occurs:

5 **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed
6 successfully - the *cbfunc* will *not* be called.

7 If none of the above return codes are appropriate, then an implementation must return either a
8 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

9 If executed, the status returned in the provided callback function will be one of the following
10 constants:

- 11 • **PMIX_SUCCESS** The operation succeeded - i.e., the **PMIX_GROUP_LEFT** event was generated.
- 12 • **PMIX_ERR_NOT_SUPPORTED** While the PMIx library supports this operation, the host RM
13 does not.
- 14 • a non-zero PMIx error constant indicating a reason for the request's failure.

Required Attributes

15 There are no identified required attributes for implementers.

Description

16 Non-blocking version of the **PMIx_Group_leave** operation. The callback function will be
17 called once the event has been locally generated and is not indicative of remote receipt.
18

CHAPTER 15

Fabric Support Definitions

1 As the drive for performance continues, interest has grown in scheduling algorithms that take into
2 account network locality of the allocated resources and in optimizing collective communication
3 patterns by structuring them to follow fabric topology. In addition, concerns over the time required
4 to initiate execution of parallel applications and enable communication across them have grown as
5 the size of those applications extends into the hundreds of thousands of individual processes
6 spanning tens of thousands of nodes.

7 PMIx supports the communication part of these efforts by defining data types and attributes by
8 which fabric endpoints and coordinates for processes and devices can be obtained from the host
9 environment. When used in conjunction with other PMIx methods described in Chapter 17, this
10 results in the ability of a process to obtain the fabric endpoint and coordinate of all other processes
11 without incurring additional overhead associated with a global exchange of that information. This
12 includes:

- 13 • Defining several interfaces specifically intended to support WLMs by providing access to
14 information of potential use to scheduling algorithms - e.g., information on communication costs
15 between different points on the fabric.
- 16 • Supporting hierarchical collective operations by providing the fabric coordinates for all devices
17 on participating nodes as well as a list of the peers sharing each fabric switch. This enables one,
18 for example, to aggregate the contribution from all processes on a node, then again across all
19 nodes on a common switch, and finally across all switches based on detailed knowledge of the
20 fabric location of each participant.
- 21 • Enabling the "*instant on*" paradigm to mitigate the scalable launch problem by providing each
22 process with a rich set of information about the environment and the application, including
23 everything required for communication between peers within the application, at time of process
24 start of execution.

25 Meeting these needs in the case where only a single fabric device exists on each node is relatively
26 straightforward - PMIx and the host environment provide a single endpoint for each process plus a
27 coordinate for the device on each node, and there is no uncertainty regarding the endpoint each
28 process will use. Extending this to the multiple device per node case is more difficult as the choice
29 of endpoint by any given process cannot be known in advance, and questions arise regarding
30 reachability between devices on different nodes. Resolving these ambiguities without requiring a
31 global operation requires that PMIx provide both (a) an endpoint for each application process on
32 each of its local devices; and (b) the fabric coordinates of all remote and local devices on
33 participating nodes. It also requires that each process open all of its assigned endpoints as the
34 endpoint selected for contact by a remote peer cannot be known in advance.

1 While these steps ensure the ability of a process to connect to a remote peer, it leaves unanswered
2 the question of selecting the *preferred* device for that communication. If multiple devices are
3 present on a node, then the application can benefit from having each process utilize its "closest"
4 fabric device (i.e., the device that minimizes the communication distance between the process'
5 location and that device) for messaging operations. In some cases, messaging libraries prefer to
6 also retain the ability to use non-nearest devices, prioritizing the devices based on distance to
7 support multi-device operations (e.g., for large message transmission in parallel).

8 PMIx supports this requirement by providing the array of process-to-device distance information
9 for each process and local fabric device at start of execution. Both minimum and maximum
10 distances are provided since a single process can occupy multiple processor locations. In addition,
11 since processes can relocate themselves by changing their processor bindings, PMIx provides an
12 API that allows the process to dynamically request an update to its distance array.

13 However, while these measures assist a process in selecting its own best endpoint, they do not
14 resolve the uncertainty over the choice of preferred device by a remote peer. There are two methods
15 by which this ambiguity can be resolved:

- 16 a) A process can select a remote endpoint to use based on its own preferred device and reachability
17 of the peer's remote devices. Once the initial connection has been made, the two processes can
18 exchange information and mutually determine their desired communication path going forward.
- 19 b) The application can use knowledge of both the local and remote distance arrays to compute the
20 best communication path and establish that connection. In some instances (e.g., a homogeneous
21 system), a PMIx server may provide distance information for both local and remote devices.
22 Alternatively, when this isn't available, an application can opt to collect the information using
23 the `PMIX_COLLECT_GENERATED_JOB_INFO` with the `PMIx_Fence` API, or can obtain it
24 on a one peer-at-a-time basis using the `PMIx_Get` API on systems where the host environment
25 supports the *Direct Modex* operation.

26 Information on fabric coordinates, endpoints, and device distances are provided as *reserved keys* as
27 detailed in Chapter 6 - i.e., they are to be available at client start of execution and are subject to the
28 retrieval rules of Section 6.2. Examples for retrieving fabric-related information include retrieval of:

- 29 • An array of information on fabric devices for a node by passing `PMIX_FABRIC_DEVICES` as
30 the key to `PMIx_Get` along with the `PMIX_HOSTNAME` of the node as a directive
- 31 • An array of information on a specific fabric device by passing `PMIX_FABRIC_DEVICE` as the
32 key to `PMIx_Get` along with the `PMIX_DEVICE_ID` of the device as a directive
- 33 • An array of information on a specific fabric device by passing `PMIX_FABRIC_DEVICE` as the
34 key to `PMIx_Get` along with both `PMIX_FABRIC_DEVICE_NAME` of the device and the
35 `PMIX_HOSTNAME` of the node as directives

36 When requesting data on a device, returned data must include at least the following attributes:

- 37 • `PMIX_HOSTNAME` "pmix.hname" (char*)

1 Name of the host, as returned by the `gethostname` utility or its equivalent. The
2 `PMIX_NODEID` may be returned in its place, or in addition to the hostname.

- 3 ● `PMIX_DEVICE_ID` `"pmix.dev.id"` (string)
4 System-wide UUID or node-local OS name of a particular device.
- 5 ● `PMIX_FABRIC_DEVICE_NAME` `"pmix.fabdev.nm"` (string)
6 The operating system name associated with the device. This may be a logical fabric
7 interface name (e.g. "eth0" or "eno1") or an absolute filename.
- 8 ● `PMIX_FABRIC_DEVICE_VENDOR` `"pmix.fabdev.vndr"` (string)
9 Indicates the name of the vendor that distributes the device.
- 10 ● `PMIX_FABRIC_DEVICE_BUS_TYPE` `"pmix.fabdev.btyp"` (string)
11 The type of bus to which the device is attached (e.g., "PCI", "GEN-Z").
- 12 ● `PMIX_FABRIC_DEVICE_PCI_DEVID` `"pmix.fabdev.pcidevid"` (string)
13 A node-level unique identifier for a Peripheral Component Interconnect (PCI) device.
14 Provided only if the device is located on a PCI bus. The identifier is constructed as a
15 four-part tuple delimited by colons comprised of the PCI 16-bit domain, 8-bit bus, 8-bit
16 device, and 8-bit function IDs, each expressed in zero-extended hexadecimal form. Thus,
17 an example identifier might be "abc1:0f:23:01". The combination of node identifier
18 (`PMIX_HOSTNAME` or `PMIX_NODEID`) and `PMIX_FABRIC_DEVICE_PCI_DEVID`
19 shall be unique within the overall system. This item should be included if the device bus
20 type is PCI - the equivalent should be provided for any other bus type.

21 The returned array may optionally contain one or more of the following in addition to the above list:

- 22 ● `PMIX_FABRIC_DEVICE_INDEX` `"pmix.fabdev.idx"` (uint32_t)
23 Index of the device within an associated communication cost matrix.
- 24 ● `PMIX_FABRIC_DEVICE_VENDORID` `"pmix.fabdev.vendid"` (string)
25 This is a vendor-provided identifier for the device or product.
- 26 ● `PMIX_FABRIC_DEVICE_DRIVER` `"pmix.fabdev.driver"` (string)
27 The name of the driver associated with the device.
- 28 ● `PMIX_FABRIC_DEVICE_FIRMWARE` `"pmix.fabdev.fmwr"` (string)
29 The device's firmware version.
- 30 ● `PMIX_FABRIC_DEVICE_ADDRESS` `"pmix.fabdev.addr"` (string)
31 The primary link-level address associated with the device, such as a Media Access
32 Control (MAC) address. If multiple addresses are available, only one will be reported.
- 33 ● `PMIX_FABRIC_DEVICE_COORDINATES` `"pmix.fab.coord"` (`pmix_geometry_t`)
34 The `pmix_geometry_t` fabric coordinates for the device, including values for all
35 supported coordinate views.

- 1 • **PMIX_FABRIC_DEVICE_MTU** "pmix.fabdev.mtu" (**size_t**)
2 The maximum transfer unit of link level frames or packets, in bytes.
- 3 • **PMIX_FABRIC_DEVICE_SPEED** "pmix.fabdev.speed" (**size_t**)
4 The active link data rate, given in bits per second.
- 5 • **PMIX_FABRIC_DEVICE_STATE** "pmix.fabdev.state" (**pmix_link_state_t**)
6 The last available physical port state for the specified device. Possible values are
7 **PMIX_LINK_STATE_UNKNOWN**, **PMIX_LINK_DOWN**, and **PMIX_LINK_UP**, to
8 indicate if the port state is unknown or not applicable (unknown), inactive (down), or
9 active (up).
- 10 • **PMIX_FABRIC_DEVICE_TYPE** "pmix.fabdev.type" (**string**)
11 Specifies the type of fabric interface currently active on the device, such as Ethernet or
12 InfiniBand.

13 The remainder of this chapter details the events, data types, attributes, and APIs associated with
14 fabric-related operations.

15 15.1 Fabric Support Events

16 The following events are defined for use in fabric-related operations.

17 **PMIX_FABRIC_UPDATE_PENDING** -176 The PMIx server library has been alerted to a
18 change in the fabric that requires updating of one or more registered **pmix_fabric_t**
19 objects.

20 **PMIX_FABRIC_UPDATED** -175 The PMIx server library has completed updating the
21 entries of all affected **pmix_fabric_t** objects registered with the library. Access to the
22 entries of those objects may now resume.

23 **PMIX_FABRIC_UPDATE_ENDPOINTS** -113 Endpoint assignments have been updated,
24 usually in response to migration or restart of a process. Clients should use **PMIx_Get** to
25 update any internally cached connections.

26 15.2 Fabric Support Datatypes

27 Several datatype definitions have been created to support fabric-related operations and information.

28 15.2.1 Fabric Endpoint Structure

29 The **pmix_endpoint_t** structure contains an assigned endpoint for a given fabric device.

PMIx v4.0

```

1 typedef struct pmix_endpoint {
2     char *uuid;
3     char *osname;
4     pmix_byte_object_t endpt;
5 } pmix_endpoint_t;

```

The *uuid* field contains the UUID of the fabric device, the *osname* is the local operating system's name for the device, and the *endpt* field contains a fabric vendor-specific object identifying the communication endpoint assigned to the process.

15.2.2 Fabric endpoint support macros

The following macros are provided to support the `pmix_endpoint_t` structure.

Static initializer for the endpoint structure

Provisional

Provide a static initializer for the `pmix_endpoint_t` fields.

PMIx v4.2

```

PMIX_ENDPOINT_STATIC_INIT

```

Initialize the endpoint structure

Initialize the `pmix_endpoint_t` fields.

PMIx v4.0

```

PMIX_ENDPOINT_CONSTRUCT (m)

```

IN m

Pointer to the structure to be initialized (pointer to `pmix_endpoint_t`)

Destruct the endpoint structure

Destruct the `pmix_endpoint_t` fields.

PMIx v4.0

```

PMIX_ENDPOINT_DESTRUCT (m)

```

IN m

Pointer to the structure to be destructed (pointer to `pmix_endpoint_t`)

1 **Create an endpoint array**
2 Allocate and initialize a `pmix_endpoint_t` array.

▼ `C` ————— ▼

3 **PMIX_ENDPOINT_CREATE**(m, n)

▲ `C` ————— ▲

4 **INOUT** m
5 Address where the pointer to the array of `pmix_endpoint_t` structures shall be stored
6 (handle)

7 **IN** n
8 Number of structures to be allocated (`size_t`)

9 **Release an endpoint array**
10 Release an array of `pmix_endpoint_t` structures.

PMIx v4.0 ▼ `C` ————— ▼

11 **PMIX_ENDPOINT_FREE**(m, n)

▲ `C` ————— ▲

12 **IN** m
13 Pointer to the array of `pmix_endpoint_t` structures (handle)

14 **IN** n
15 Number of structures in the array (`size_t`)

16 15.2.3 Fabric Coordinate Structure

17 The `pmix_coord_t` structure describes the fabric coordinates of a specified device in a given
18 view.

PMIx v4.0 ▼ `C` ————— ▼

```
19 typedef struct pmix_coord {  
20     pmix_coord_view_t view;  
21     uint32_t *coord;  
22     size_t dims;  
23 } pmix_coord_t;
```

▲ `C` ————— ▲

24 All coordinate values shall be expressed as unsigned integers due to their units being defined in
25 fabric devices and not physical distances. The coordinate is therefore an indicator of connectivity
26 and not relative communication distance.

▼ **Advice to PMIx library implementers** ————— ▼

27 Note that the `pmix_coord_t` structure does not imply nor mandate any requirement on how the
28 coordinate data is to be stored within the PMIx library. Implementers are free to store the
29 coordinate in whatever format they choose.

1 A fabric coordinate is associated with a given fabric device and must be unique within a given view.
2 Fabric devices are associated with the operating system which hosts them - thus, fabric coordinates
3 are logically grouped within the *node* realm (as described in Section 6.1) and can be retrieved per
4 the rules detailed in Section 6.1.5.

5 15.2.4 Fabric coordinate support macros

6 The following macros are provided to support the `pmix_coord_t` structure.

7 **Static initializer for the coord structure**

8 *Provisional*

Provide a static initializer for the `pmix_coord_t` fields.

PMIx v4.2

9 **PMIX_COORD_STATIC_INIT**

10 **Initialize the coord structure**

11 Initialize the `pmix_coord_t` fields.

PMIx v4.0

12 **PMIX_COORD_CONSTRUCT (m)**

13 **IN** m

14 Pointer to the structure to be initialized (pointer to `pmix_coord_t`)

15 **Destruct the coord structure**

16 Destruct the `pmix_coord_t` fields.

PMIx v4.0

17 **PMIX_COORD_DESTRUCT (m)**

18 **IN** m

19 Pointer to the structure to be destructed (pointer to `pmix_coord_t`)

20 **Create a coord array**

21 Allocate and initialize a `pmix_coord_t` array.

PMIx v4.0

22 **PMIX_COORD_CREATE (m, n)**

23 **INOUT** m

24 Address where the pointer to the array of `pmix_coord_t` structures shall be stored (handle)

25 **IN** n

26 Number of structures to be allocated (`size_t`)

1 **Release a coord array**
2 Release an array of `pmix_coord_t` structures.

PMIx v4.0

▼ **C** ————— ▼

3 **PMIX_COORD_FREE**(*m*, *n*)

▲ ————— **C** ————— ▲

4 **IN** *m*
5 Pointer to the array of `pmix_coord_t` structures (handle)
6 **IN** *n*
7 Number of structures in the array (`size_t`)

8 15.2.5 Fabric Geometry Structure

9 The `pmix_geometry_t` structure describes the fabric coordinates of a specified device.

PMIx v4.0

▼ **C** ————— ▼

```
10 typedef struct pmix_geometry {  
11     size_t fabric;  
12     char *uuid;  
13     char *osname;  
14     pmix_coord_t *coordinates;  
15     size_t ncoords;  
16 } pmix_geometry_t;
```

▲ ————— **C** ————— ▲

17 All coordinate values shall be expressed as unsigned integers due to their units being defined in
18 fabric devices and not physical distances. The coordinate is therefore an indicator of connectivity
19 and not relative communication distance.

▼ **Advice to PMIx library implementers** ————— ▼

20 Note that the `pmix_coord_t` structure does not imply nor mandate any requirement on how the
21 coordinate data is to be stored within the PMIx library. Implementers are free to store the
22 coordinate in whatever format they choose.

▲ ————— ▲

23 A fabric coordinate is associated with a given fabric device and must be unique within a given view.
24 Fabric devices are associated with the operating system which hosts them - thus, fabric coordinates
25 are logically grouped within the *node* realm (as described in Section 6.1) and can be retrieved per
26 the rules detailed in Section 6.1.5.

27 15.2.6 Fabric geometry support macros

28 The following macros are provided to support the `pmix_geometry_t` structure.

1 **Static initializer for the geometry structure**

2 Provide a static initializer for the `pmix_geometry_t` fields.



3 `PMIX_GEOMETRY_STATIC_INIT`



4 **Initialize the geometry structure**

5 Initialize the `pmix_geometry_t` fields.

PMIx v4.0



6 `PMIX_GEOMETRY_CONSTRUCT (m)`



7 **IN** `m`

8 Pointer to the structure to be initialized (pointer to `pmix_geometry_t`)

9 **Destruct the geometry structure**

10 Destruct the `pmix_geometry_t` fields.

PMIx v4.0



11 `PMIX_GEOMETRY_DESTRUCT (m)`



12 **IN** `m`

13 Pointer to the structure to be destructed (pointer to `pmix_geometry_t`)

14 **Create a geometry array**

15 Allocate and initialize a `pmix_geometry_t` array.

PMIx v4.0



16 `PMIX_GEOMETRY_CREATE (m, n)`



17 **INOUT** `m`

18 Address where the pointer to the array of `pmix_geometry_t` structures shall be stored
19 (handle)

20 **IN** `n`

21 Number of structures to be allocated (`size_t`)

1 **Release a geometry array**
2 Release an array of `pmix_geometry_t` structures.

3 `PMIX_GEOMETRY_FREE(m, n)` C

4 **PMIX_GEOMETRY_FREE(m, n)**

5 `IN m` C

6 `IN n`
7 Pointer to the array of `pmix_geometry_t` structures (handle)

8 `IN n`
9 Number of structures in the array (`size_t`)

8 15.2.7 Fabric Coordinate Views

PMIx v4.0

C

```
9 typedef uint8_t pmix_coord_view_t;  
10 #define PMIX_COORD_VIEW_UNDEF 0x00  
11 #define PMIX_COORD_LOGICAL_VIEW 0x01  
12 #define PMIX_COORD_PHYSICAL_VIEW 0x02
```

C

13 Fabric coordinates can be reported based on different *views* according to user preference at the time
14 of request. The following views have been defined:

15 **PMIX_COORD_VIEW_UNDEF 0x00** The coordinate view has not been defined.

16 **PMIX_COORD_LOGICAL_VIEW 0x01** The coordinates are provided in a *logical* view,
17 typically given in Cartesian (x,y,z) dimensions, that describes the data flow in the fabric as
18 defined by the arrangement of the hierarchical addressing scheme, fabric segmentation,
19 routing domains, and other similar factors employed by that fabric.

20 **PMIX_COORD_PHYSICAL_VIEW 0x02** The coordinates are provided in a *physical* view
21 based on the actual wiring diagram of the fabric - i.e., values along each axis reflect the
22 relative position of that interface on the specific fabric cabling.

23 If the requester does not specify a view, coordinates shall default to the *logical* view.

24 15.2.8 Fabric Link State

25 The `pmix_link_state_t` is a `uint32_t` type for fabric link states.

PMIx v4.0

```
1 typedef uint8_t pmix_link_state_t;
```

2 The following constants can be used to set a variable of the type `pmix_link_state_t`. All
3 definitions were introduced in version 4 of the standard unless otherwise marked. Valid link state
4 values start at zero.

5 **PMIX_LINK_STATE_UNKNOWN** 0 The port state is unknown or not applicable.

6 **PMIX_LINK_DOWN** 1 The port is inactive.

7 **PMIX_LINK_UP** 2 The port is active.

8 15.2.9 Fabric Operation Constants

9 *PMIx v4.0*

10 The `pmix_fabric_operation_t` data type is an enumerated type for specifying fabric
11 operations used in the PMIx server module's `pmix_server_fabric_fn_t` API.

12 **PMIX_FABRIC_REQUEST_INFO** 0 Request information on a specific fabric - if the fabric
13 isn't specified as per `PMIx_Fabric_register`, then return information on the default
14 fabric of the overall system. Information to be returned is described in `pmix_fabric_t`.

15 **PMIX_FABRIC_UPDATE_INFO** 1 Update information on a specific fabric - the index of the
16 fabric (`PMIX_FABRIC_INDEX`) to be updated must be provided.

16 15.2.10 Fabric registration structure

17 The `pmix_fabric_t` structure is used by a WLM to interact with fabric-related PMIx interfaces,
18 and to provide information about the fabric for use in scheduling algorithms or other purposes.

19 *PMIx v4.0*

```
19 typedef struct pmix_fabric_s {  
20     char *name;  
21     size_t index;  
22     pmix_info_t *info;  
23     size_t ninfo;  
24     void *module;  
25 } pmix_fabric_t;
```

26 Note that in this structure:

- 27 • *name* is an optional user-supplied string name identifying the fabric being referenced by this
28 struct. If provided, the field must be a **NULL**-terminated string composed of standard
29 alphanumeric values supported by common utilities such as *strcmp*;
- 30 • *index* is a PMIx-provided number identifying this object;

- 1 • *info* is an array of `pmix_info_t` containing information (provided by the PMIx library) about
2 the fabric;
- 3 • *ninfo* is the number of elements in the *info* array;
- 4 • *module* points to an opaque object reserved for use by the PMIx server library.

5 Note that only the *name* field is provided by the user - all other fields are provided by the PMIx
6 library and must not be modified by the user. The *info* array contains a varying amount of
7 information depending upon both the PMIx implementation and information available from the
8 fabric vendor. At a minimum, it must contain (ordering is arbitrary):

Required Attributes

- 9 **PMIX_FABRIC_VENDOR** "`pmix.fab.vndr`" (**string**)
10 Name of the vendor (e.g., Amazon, Mellanox, HPE, Intel) for the specified fabric.
- 11 **PMIX_FABRIC_IDENTIFIER** "`pmix.fab.id`" (**string**)
12 An identifier for the specified fabric (e.g., MgmtEthernet, Slingshot-11, OmniPath-1).
- 13 **PMIX_FABRIC_NUM_DEVICES** "`pmix.fab.nverts`" (**size_t**)
14 Total number of fabric devices in the overall system - corresponds to the number of rows or
15 columns in the cost matrix.

16 and may optionally contain one or more of the following:

Optional Attributes

- 17 **PMIX_FABRIC_COST_MATRIX** "`pmix.fab.cm`" (**pointer**)
18 Pointer to a two-dimensional square array of point-to-point relative communication costs
19 expressed as `uint16_t` values.
- 20 **PMIX_FABRIC_GROUPS** "`pmix.fab.grps`" (**string**)
21 A string delineating the group membership of nodes in the overall system, where each fabric
22 group consists of the group number followed by a colon and a comma-delimited list of nodes
23 in that group, with the groups delimited by semi-colons (e.g.,
24 `0:node000,node002,node004,node006;1:node001,node003,`
25 `node005,node007`)
- 26 **PMIX_FABRIC_DIMS** "`pmix.fab.dims`" (**uint32_t**)
27 Number of dimensions in the specified fabric plane/view. If no plane is specified in a
28 request, then the dimensions of all planes in the overall system will be returned as a
29 `pmix_data_array_t` containing an array of `uint32_t` values. Default is to provide
30 dimensions in *logical* view.
- 31 **PMIX_FABRIC_PLANE** "`pmix.fab.plane`" (**string**)

1 ID string of a fabric plane (e.g., CIDR for Ethernet). When used as a modifier in a request
2 for information, specifies the plane whose information is to be returned. When used directly
3 as a key in a request, returns a `pmix_data_array_t` of string identifiers for all fabric
4 planes in the overall system.

5 **PMIX_FABRIC_SHAPE** "`pmix.fab.shape`" (`pmix_data_array_t*`)

6 The size of each dimension in the specified fabric plane/view, returned in a
7 `pmix_data_array_t` containing an array of `uint32_t` values. The size is defined as
8 the number of elements present in that dimension - e.g., the number of devices in one
9 dimension of a physical view of a fabric plane. If no plane is specified, then the shape of
10 each plane in the overall system will be returned in a `pmix_data_array_t` array where
11 each element is itself a two-element array containing the **PMIX_FABRIC_PLANE** followed
12 by that plane's fabric shape. Default is to provide the shape in *logical* view.

13 **PMIX_FABRIC_SHAPE_STRING** "`pmix.fab.shapestr`" (`string`)

14 Network shape expressed as a string (e.g., "`10x12x2`"). If no plane is specified, then the
15 shape of each plane in the overall system will be returned in a `pmix_data_array_t` array
16 where each element is itself a two-element array containing the **PMIX_FABRIC_PLANE**
17 followed by that plane's fabric shape string. Default is to provide the shape in *logical* view.

18 While unusual due to scaling issues, implementations may include an array of
19 **PMIX_FABRIC_DEVICE** elements describing the device information for each device in the
20 overall system. Each element shall contain a `pmix_data_array_t` of `pmix_info_t` values
21 describing the device. Each array may contain one or more of the following (ordering is arbitrary):

22 **PMIX_FABRIC_DEVICE_NAME** "`pmix.fabdev.nm`" (`string`)

23 The operating system name associated with the device. This may be a logical fabric interface
24 name (e.g. "`eth0`" or "`eno1`") or an absolute filename.

25 **PMIX_FABRIC_DEVICE_VENDOR** "`pmix.fabdev.vndr`" (`string`)

26 Indicates the name of the vendor that distributes the device.

27 **PMIX_DEVICE_ID** "`pmix.dev.id`" (`string`)

28 System-wide UUID or node-local OS name of a particular device.

29 **PMIX_HOSTNAME** "`pmix.hname`" (`char*`)

30 Name of the host, as returned by the `gethostname` utility or its equivalent.

31 **PMIX_FABRIC_DEVICE_DRIVER** "`pmix.fabdev.driver`" (`string`)

32 The name of the driver associated with the device.

33 **PMIX_FABRIC_DEVICE_FIRMWARE** "`pmix.fabdev.fmwr`" (`string`)

34 The device's firmware version.

35 **PMIX_FABRIC_DEVICE_ADDRESS** "`pmix.fabdev.addr`" (`string`)

36 The primary link-level address associated with the device, such as a MAC address. If
37 multiple addresses are available, only one will be reported.

38 **PMIX_FABRIC_DEVICE_MTU** "`pmix.fabdev.mtu`" (`size_t`)

The maximum transfer unit of link level frames or packets, in bytes.

PMIX_FABRIC_DEVICE_SPEED "pmix.fabdev.speed" (**size_t**)

The active link data rate, given in bits per second.

PMIX_FABRIC_DEVICE_STATE "pmix.fabdev.state" (**pmix_link_state_t**)

The last available physical port state for the specified device. Possible values are **PMIX_LINK_STATE_UNKNOWN**, **PMIX_LINK_DOWN**, and **PMIX_LINK_UP**, to indicate if the port state is unknown or not applicable (unknown), inactive (down), or active (up).

PMIX_FABRIC_DEVICE_TYPE "pmix.fabdev.type" (**string**)

Specifies the type of fabric interface currently active on the device, such as Ethernet or InfiniBand.

PMIX_FABRIC_DEVICE_BUS_TYPE "pmix.fabdev.btyp" (**string**)

The type of bus to which the device is attached (e.g., "PCI", "GEN-Z").

PMIX_FABRIC_DEVICE_PCI_DEVID "pmix.fabdev.pcidevid" (**string**)

A node-level unique identifier for a PCI device. Provided only if the device is located on a PCI bus. The identifier is constructed as a four-part tuple delimited by colons comprised of the PCI 16-bit domain, 8-bit bus, 8-bit device, and 8-bit function IDs, each expressed in zero-extended hexadecimal form. Thus, an example identifier might be "abc1:0f:23:01". The combination of node identifier (**PMIX_HOSTNAME** or **PMIX_NODEID**) and **PMIX_FABRIC_DEVICE_PCI_DEVID** shall be unique within the overall system.

15.2.10.1 Static initializer for the fabric structure

Provisional

Provide a static initializer for the **pmix_fabric_t** fields.

PMIx v4.2

PMIX_FABRIC_STATIC_INIT

C

15.2.10.2 Initialize the fabric structure

Initialize the **pmix_fabric_t** fields.

PMIx v4.0

PMIX_FABRIC_CONSTRUCT (m)

C

IN m

Pointer to the structure to be initialized (pointer to **pmix_fabric_t**)

1 15.3 Fabric Support Attributes

2 The following attribute is used by the PMIx server library supporting the system's WLM to indicate
3 that it wants access to the fabric support functions:

4 **PMIX_SERVER_SCHEDULER** "pmix.srv.sched" (bool)

5 Server is supporting system scheduler and desires access to appropriate WLM-supporting
6 features. Indicates that the library is to be initialized for scheduler support.

7 The following attributes may be returned in response to fabric-specific APIs or queries (e.g.,
8 **PMIx_Get** or **PMIx_Query_info**). These attributes are not related to a specific *data realm* (as
9 described in Section 6.1) - the **PMIx_Get** function shall therefore ignore the value in its *proc*
10 process identifier argument when retrieving these values.

11 **PMIX_FABRIC_COST_MATRIX** "pmix.fab.cm" (pointer)

12 Pointer to a two-dimensional square array of point-to-point relative communication costs
13 expressed as `uint16_t` values.

14 **PMIX_FABRIC_GROUPS** "pmix.fab.grps" (string)

15 A string delineating the group membership of nodes in the overall system, where each fabric
16 group consists of the group number followed by a colon and a comma-delimited list of nodes
17 in that group, with the groups delimited by semi-colons (e.g.,
18 0:node000,node002,node004,node006;1:node001,node003,
19 node005,node007)

20 **PMIX_FABRIC_PLANE** "pmix.fab.plane" (string)

21 ID string of a fabric plane (e.g., CIDR for Ethernet). When used as a modifier in a request
22 for information, specifies the plane whose information is to be returned. When used directly
23 as a key in a request, returns a `pmix_data_array_t` of string identifiers for all fabric
24 planes in the overall system.

25 **PMIX_FABRIC_SWITCH** "pmix.fab.switch" (string)

26 ID string of a fabric switch. When used as a modifier in a request for information, specifies
27 the switch whose information is to be returned. When used directly as a key in a request,
28 returns a `pmix_data_array_t` of string identifiers for all fabric switches in the overall
29 system.

30 The following attributes may be returned in response to queries (e.g., **PMIx_Get** or
31 **PMIx_Query_info**). A qualifier (e.g., **PMIX_FABRIC_INDEX**) identifying the fabric whose
32 value is being referenced must be provided for queries on systems supporting more than one fabric
33 when values for the non-default fabric are requested. These attributes are not related to a specific
34 *data realm* (as described in Section 6.1) - the **PMIx_Get** function shall therefore ignore the value
35 in its *proc* process identifier argument when retrieving these values.

36 **PMIX_FABRIC_VENDOR** "pmix.fab.vndr" (string)

37 Name of the vendor (e.g., Amazon, Mellanox, HPE, Intel) for the specified fabric.

38 **PMIX_FABRIC_IDENTIFIER** "pmix.fab.id" (string)

1 An identifier for the specified fabric (e.g., MgmtEthernet, Slingshot-11, OmniPath-1).

2 **PMIX_FABRIC_INDEX** "pmix.fab.idx" (**size_t**)

3 The index of the fabric as returned in [pmix_fabric_t](#).

4 **PMIX_FABRIC_NUM_DEVICES** "pmix.fab.nverts" (**size_t**)

5 Total number of fabric devices in the overall system - corresponds to the number of rows or
6 columns in the cost matrix.

7 **PMIX_FABRIC_DIMS** "pmix.fab.dims" (**uint32_t**)

8 Number of dimensions in the specified fabric plane/view. If no plane is specified in a
9 request, then the dimensions of all planes in the overall system will be returned as a
10 [pmix_data_array_t](#) containing an array of **uint32_t** values. Default is to provide
11 dimensions in *logical* view.

12 **PMIX_FABRIC_SHAPE** "pmix.fab.shape" (**pmix_data_array_t***)

13 The size of each dimension in the specified fabric plane/view, returned in a
14 [pmix_data_array_t](#) containing an array of **uint32_t** values. The size is defined as
15 the number of elements present in that dimension - e.g., the number of devices in one
16 dimension of a physical view of a fabric plane. If no plane is specified, then the shape of
17 each plane in the overall system will be returned in a [pmix_data_array_t](#) array where
18 each element is itself a two-element array containing the [PMIX_FABRIC_PLANE](#) followed
19 by that plane's fabric shape. Default is to provide the shape in *logical* view.

20 **PMIX_FABRIC_SHAPE_STRING** "pmix.fab.shapeststr" (**string**)

21 Network shape expressed as a string (e.g., "10x12x2"). If no plane is specified, then the
22 shape of each plane in the overall system will be returned in a [pmix_data_array_t](#) array
23 where each element is itself a two-element array containing the [PMIX_FABRIC_PLANE](#)
24 followed by that plane's fabric shape string. Default is to provide the shape in *logical* view.

25 The following attributes are related to the *node realm* (as described in Section 6.1.5) and are
26 retrieved according to those rules.

27 **PMIX_FABRIC_DEVICES** "pmix.fab.devs" (**pmix_data_array_t**)

28 Array of [pmix_info_t](#) containing information for all devices on the specified node. Each
29 element of the array will contain a [PMIX_FABRIC_DEVICE](#) entry, which in turn will
30 contain an array of information on a given device.

31 **PMIX_FABRIC_COORDINATES** "pmix.fab.coords" (**pmix_data_array_t**)

32 Array of [pmix_geometry_t](#) fabric coordinates for devices on the specified node. The
33 array will contain the coordinates of all devices on the node, including values for all
34 supported coordinate views. The information for devices on the local node shall be provided
35 if the node is not specified in the request.

36 **PMIX_FABRIC_DEVICE** "pmix.fabdev" (**pmix_data_array_t**)

37 An array of [pmix_info_t](#) describing a particular fabric device using one or more of the
38 attributes defined below. The first element in the array shall be the [PMIX_DEVICE_ID](#) of
39 the device.

40 **PMIX_FABRIC_DEVICE_INDEX** "pmix.fabdev.idx" (**uint32_t**)

41 Index of the device within an associated communication cost matrix.

42 **PMIX_FABRIC_DEVICE_NAME** "pmix.fabdev.nm" (**string**)

1 The operating system name associated with the device. This may be a logical fabric interface
2 name (e.g. "eth0" or "eno1") or an absolute filename.

3 **PMIX_FABRIC_DEVICE_VENDOR** "pmix.fabdev.vndr" (string)

4 Indicates the name of the vendor that distributes the device.

5 **PMIX_FABRIC_DEVICE_BUS_TYPE** "pmix.fabdev.btyp" (string)

6 The type of bus to which the device is attached (e.g., "PCI", "GEN-Z").

7 **PMIX_FABRIC_DEVICE_VENDORID** "pmix.fabdev.vendid" (string)

8 This is a vendor-provided identifier for the device or product.

9 **PMIX_FABRIC_DEVICE_DRIVER** "pmix.fabdev.driver" (string)

10 The name of the driver associated with the device.

11 **PMIX_FABRIC_DEVICE_FIRMWARE** "pmix.fabdev.fmwr" (string)

12 The device's firmware version.

13 **PMIX_FABRIC_DEVICE_ADDRESS** "pmix.fabdev.addr" (string)

14 The primary link-level address associated with the device, such as a MAC address. If
15 multiple addresses are available, only one will be reported.

16 **PMIX_FABRIC_DEVICE_COORDINATES** "pmix.fab.coord" (pmix_geometry_t)

17 The [pmix_geometry_t](#) fabric coordinates for the device, including values for all
18 supported coordinate views.

19 **PMIX_FABRIC_DEVICE_MTU** "pmix.fabdev.mtu" (size_t)

20 The maximum transfer unit of link level frames or packets, in bytes.

21 **PMIX_FABRIC_DEVICE_SPEED** "pmix.fabdev.speed" (size_t)

22 The active link data rate, given in bits per second.

23 **PMIX_FABRIC_DEVICE_STATE** "pmix.fabdev.state" (pmix_link_state_t)

24 The last available physical port state for the specified device. Possible values are
25 [PMIX_LINK_STATE_UNKNOWN](#), [PMIX_LINK_DOWN](#), and [PMIX_LINK_UP](#), to indicate
26 if the port state is unknown or not applicable (unknown), inactive (down), or active (up).

27 **PMIX_FABRIC_DEVICE_TYPE** "pmix.fabdev.type" (string)

28 Specifies the type of fabric interface currently active on the device, such as Ethernet or
29 InfiniBand.

30 **PMIX_FABRIC_DEVICE_PCI_DEVID** "pmix.fabdev.pcidevid" (string)

31 A node-level unique identifier for a PCI device. Provided only if the device is located on a
32 PCI bus. The identifier is constructed as a four-part tuple delimited by colons comprised of
33 the PCI 16-bit domain, 8-bit bus, 8-bit device, and 8-bit function IDs, each expressed in
34 zero-extended hexadecimal form. Thus, an example identifier might be "abc1:0f:23:01". The
35 combination of node identifier ([PMIX_HOSTNAME](#) or [PMIX_NODEID](#)) and
36 [PMIX_FABRIC_DEVICE_PCI_DEVID](#) shall be unique within the overall system.

37 The following attributes are related to the *process realm* (as described in Section 6.1.4) and are
38 retrieved according to those rules.

39 **PMIX_FABRIC_ENDPT** "pmix.fab.endpt" (pmix_data_array_t)

40 Fabric endpoints for a specified process. As multiple endpoints may be assigned to a given
41 process (e.g., in the case where multiple devices are associated with a package to which the

1 process is bound), the returned values will be provided in a `pmix_data_array_t` of
2 `pmix_endpoint_t` elements.

3 The following attributes are related to the *job realm* (as described in Section 6.1.2) and are retrieved
4 according to those rules. Note that distances to fabric devices are retrieved using the
5 `PMIX_DEVICE_DISTANCES` key with the appropriate `pmix_device_type_t` qualifier.

6 **PMIX_SWITCH_PEERS** "`pmix.speers`" (`pmix_data_array_t`)

7 Peer ranks that share the same switch as the process specified in the call to `PMIx_Get`.

8 Returns a `pmix_data_array_t` array of `pmix_info_t` results, each element
9 containing the `PMIX_SWITCH_PEERS` key with a three-element `pmix_data_array_t`
10 array of `pmix_info_t` containing the `PMIX_DEVICE_ID` of the local fabric device, the
11 `PMIX_FABRIC_SWITCH` identifying the switch to which it is connected, and a
12 comma-delimited string of peer ranks sharing the switch to which that device is connected.

13 15.4 Fabric Support Functions

14 The following APIs allow the WLM to request specific services from the fabric subsystem via the
15 PMIx library.

▼ Advice to PMIx server hosts ▼

16 Due to their high cost in terms of execution, memory consumption, and interactions with other
17 SMS components (e.g., a fabric manager), it is strongly advised that the underlying implementation
18 of these APIs be restricted to a single PMIx server in a system that is supporting the SMS
19 component responsible for the scheduling of allocations (i.e., the system *scheduler*). The
20 `PMIX_SERVER_SCHEDULER` attribute can be used for this purpose to control the execution path.
21 Clients, tools, and other servers utilizing these functions are advised to have their requests
22 forwarded to the server supporting the scheduler using the `pmix_server_fabric_fn_t`
23 server module function, as needed.

24 15.4.1 PMIx_Fabric_register

25 Summary

26 Register for access to fabric-related information.

Format

```
pmix_status_t
PMIx_Fabric_register(pmix_fabric_t *fabric,
                    const pmix_info_t directives[],
                    size_t ndirs);
```

INOUT fabric

address of a `pmix_fabric_t` (backed by storage). User may populate the "name" field at will - PMIx does not utilize this field (handle)

IN directives

an optional array of values indicating desired behaviors and/or fabric to be accessed. If `NULL`, then the highest priority available fabric will be used (array of handles)

IN ndirs

Number of elements in the *directives* array (integer)

Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

The following directives are required to be supported by all PMIx libraries to aid users in identifying the fabric whose data is being sought:

`PMIX_FABRIC_PLANE` "pmix.fab.plane" (string)

ID string of a fabric plane (e.g., CIDR for Ethernet). When used as a modifier in a request for information, specifies the plane whose information is to be returned. When used directly as a key in a request, returns a `pmix_data_array_t` of string identifiers for all fabric planes in the overall system.

`PMIX_FABRIC_IDENTIFIER` "pmix.fab.id" (string)

An identifier for the specified fabric (e.g., MgmtEthernet, Slingshot-11, OmniPath-1).

`PMIX_FABRIC_VENDOR` "pmix.fab.vndr" (string)

Name of the vendor (e.g., Amazon, Mellanox, HPE, Intel) for the specified fabric.

Description

Register for access to fabric-related information, including the communication cost matrix. This call must be made prior to requesting information from a fabric. The caller may request access to a particular fabric using the vendor, type, or identifier, or to a specific *fabric plane* via the **PMIX_FABRIC_PLANE** attribute - otherwise, information for the default fabric will be returned. Upon successful completion of the call, information will have been filled into the fields of the provided *fabric* structure.

For performance reasons, the PMIx library does not provide thread protection for accessing the information in the **pmix_fabric_t** structure. Instead, the PMIx implementation shall provide two methods for coordinating updates to the provided fabric information:

- Users may periodically poll for updates using the **PMIx_Fabric_update** API
- Users may register for **PMIX_FABRIC_UPDATE_PENDING** events indicating that an update to the cost matrix is pending. When received, users are required to terminate or pause any actions involving access to the cost matrix before returning from the event. Completion of the **PMIX_FABRIC_UPDATE_PENDING** event handler indicates to the PMIx library that the fabric object's entries are available for updating. This may include releasing and re-allocating memory as the number of vertices may have changed (e.g., due to addition or removal of one or more devices). When the update has been completed, the PMIx library will generate a **PMIX_FABRIC_UPDATED** event indicating that it is safe to begin using the updated fabric object(s).

There is no requirement that the caller exclusively use either one of these options. For example, the user may choose to both register for fabric update events, but poll for an update prior to some critical operation.

15.4.2 PMIx_Fabric_register_nb

Summary

Register for access to fabric-related information.

Format

PMIx v4.0

pmix_status_t

```
PMIx_Fabric_register_nb(pmix_fabric_t *fabric,  
                        const pmix_info_t directives[],  
                        size_t ndirs,  
                        pmix_op_cbfunc_t cbfunc, void *cbdata);
```

INOUT fabric

address of a **pmix_fabric_t** (backed by storage). User may populate the "name" field at will - PMIx does not utilize this field (handle)

1 **IN directives**
 2 an optional array of values indicating desired behaviors and/or fabric to be accessed. If **NULL**,
 3 then the highest priority available fabric will be used (array of handles)
 4 **IN ndirs**
 5 Number of elements in the *directives* array (integer)
 6 **IN cbfunc**
 7 Callback function `pmix_op_cbfunc_t` (function reference)
 8 **IN cbdata**
 9 Data to be passed to the callback function (memory reference)

10 A successful return indicates that the request is being processed and the result will be returned in
 11 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
 12 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

13 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 14 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
 15 returned *success* - the *cbfunc* will not be called

16 If none of the above return codes are appropriate, then an implementation must return either a
 17 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

18 **Description**

19 Non-blocking form of `PMIx_Fabric_register`. The caller is not allowed to access the
 20 provided `pmix_fabric_t` until the callback function has been executed, at which time the fabric
 21 information will have been loaded into the provided structure.


22 **15.4.3 PMIx_Fabric_update**

23 **Summary**

24 Update fabric-related information.

25 **Format**

PMIx v4.0



28 **INOUT fabric**

29 address of a `pmix_fabric_t` (backed by storage) (handle)

30 Returns **PMIX_SUCCESS** or a negative value indicating the error.

31 **Description**

32 Update fabric-related information. This call can be made at any time to request an update of the
 33 fabric information contained in the provided `pmix_fabric_t` object. The caller is not allowed to
 34 access the provided `pmix_fabric_t` until the call has returned. Upon successful return, the
 35 information fields in the *fabric* structure will have been updated.

1 15.4.4 PMIx_Fabric_update_nb

2 Summary

3 Update fabric-related information.

4 Format

PMIx v4.0

```
pmix_status_t  
PMIx_Fabric_update_nb(pmix_fabric_t *fabric,  
                      pmix_op_cbfunc_t cbfunc, void *cbdata);
```

8 INOUT fabric

9 address of a [pmix_fabric_t](#) (handle)

10 IN cbfunc

11 Callback function [pmix_op_cbfunc_t](#) (function reference)

12 IN cbdata

13 Data to be passed to the callback function (memory reference)

14 A successful return indicates that the request is being processed and the result will be returned in
15 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
16 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

17 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 18 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
19 returned *success* - the *cbfunc* will not be called

20 If none of the above return codes are appropriate, then an implementation must return either a
21 general **PMIx** error code or an implementation defined error code as described in Section 3.1.1.

22 Description

23 Non-blocking form of [PMIx_Fabric_update](#). The caller is not allowed to access the provided
24 [pmix_fabric_t](#) until the callback function has been executed, at which time the fields in the
25 provided *fabric* structure will have been updated.

26 15.4.5 PMIx_Fabric_deregister

27 Summary

28 Deregister a fabric object.

1 **Format** C

```
2 pmix_status_t  
3 PMIx_Fabric_deregister(pmix_fabric_t *fabric);
```

4 **IN fabric**
5 address of a **pmix_fabric_t** (handle)

6 Returns **PMIX_SUCCESS** or a negative value indicating the error.

7 **Description**

8 Deregister a fabric object, providing an opportunity for the PMIx library to cleanup any information
9 (e.g., cost matrix) associated with it. Contents of the provided **pmix_fabric_t** will be
10 invalidated upon function return.

11 **15.4.6 PMIx_Fabric_deregister_nb**

12 **Summary**

13 Deregister a fabric object.

14 *PMIx v4.0* **Format** C

```
15 pmix_status_t PMIx_Fabric_deregister_nb(pmix_fabric_t *fabric,  
16 pmix_op_cbfunc_t cbfunc,  
17 void *cbdata);
```

18 **IN fabric**
19 address of a **pmix_fabric_t** (handle)

20 **IN cbfunc**
21 Callback function **pmix_op_cbfunc_t** (function reference)

22 **IN cbdata**
23 Data to be passed to the callback function (memory reference)

24 A successful return indicates that the request is being processed and the result will be returned in
25 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
26 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

27 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 28 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
29 returned *success* - the *cbfunc* will not be called

30 If none of the above return codes are appropriate, then an implementation must return either a
31 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

32 **Description**

33 Non-blocking form of **PMIx_Fabric_deregister**. Provided *fabric* must not be accessed until
34 after callback function has been executed.

CHAPTER 16

Security

1 PMIx utilizes a multi-layered approach toward security that differs for client versus tool processes.
2 By definition, *client* processes must be preregistered with the PMIx server library via the
3 `PMIx_server_register_client` API before they are spawned. This API requires that the
4 host pass the expected effective UID/GID of the client process.

5 When the client attempts to connect to the PMIx server, the server shall use available standard OS
6 methods to determine the effective UID/GID of the process requesting the connection. PMIx
7 implementations shall not rely on any values reported by the client process itself. The effective
8 UID/GID reported by the OS is compared to the values provided by the host during registration - if
9 the values fail to match, the PMIx server is required to drop the connection request. This ensures
10 that the PMIx server does not allow connection from a client that doesn't at least meet some
11 minimal security requirement.

12 Once the requesting client passes the initial test, the PMIx server can, at the choice of the
13 implementor, perform additional security checks. This may involve a variety of methods such as
14 exchange of a system-provided key or credential. At the conclusion of that process, the PMIx server
15 reports the client connection request to the host via the
16 `pmix_server_client_connected2_fn_t` interface, if provided. The host may perform
17 any additional checks and operations before responding with either `PMIX_SUCCESS` to indicate
18 that the connection is approved, or a PMIx error constant indicating that the connection request is
19 refused. In this latter case, the PMIx server is required to drop the connection.

20 Tools started by the host environment are classed as a subgroup of client processes and follow the
21 client process procedure. However, tools that are not started by the host environment must be
22 handled differently as registration information is not available prior to the connection request. In
23 these cases, the PMIx server library is required to use available standard OS methods to get the
24 effective UID/GID of the tool and report them upwards as part of invoking the
25 `pmix_server_tool_connection_fn_t` interface, deferring initial security screening to the
26 host. Host environments willing to accept tool connections must therefore both explicitly enable
27 them via the `PMIX_SERVER_TOOL_SUPPORT` attribute, thereby confirming acceptance of the
28 authentication and authorization burden, and provide the
29 `pmix_server_tool_connection_fn_t` server module function pointer.

30 16.1 Obtaining Credentials

31 Applications and tools often interact with the host environment in ways that require security beyond
32 just verifying the user's identity - e.g., access to that user's relevant authorizations. This is

1 particularly important when tools connect directly to a system-level PMIx server that may be
2 operating at a privileged level. A variety of system management software packages provide
3 authorization services, but the lack of standardized interfaces makes portability problematic.

4 This section defines two PMIx client-side APIs for this purpose. These are most likely to be used
5 by user-space applications/tools, but are not restricted to that realm.

6 16.1.1 PMIx_Get_credential

7 Summary

8 Request a credential from the PMIx server library or the host environment.

9 Format

C

PMIx v3.0

```
10 pmix_status_t  
11 PMIx_Get_credential(const pmix_info_t info[], size_t ninfo,  
12 pmix_byte_object_t *credential);
```

C

13 IN info

14 Array of `pmix_info_t` structures (array of handles)

15 IN ninfo

16 Number of elements in the *info* array (`size_t`)

17 IN credential

18 Address of a `pmix_byte_object_t` within which to return credential (handle)

19 A successful return indicates that the credential has been returned in the provided
20 `pmix_byte_object_t`.

21 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

22 There are no required attributes for this API. Note that implementations may choose to internally
23 execute integration for some security environments (e.g., directly contacting a *munge* server).

24 Implementations that support the operation but cannot directly process the client's request must
25 pass any attributes that are provided by the client to the host environment for processing. In
26 addition, the following attributes are required to be included in the *info* array passed from the PMIx
27 library to the host environment:

28 `PMIX_USERID` "pmix.euid" (`uint32_t`)

29 Effective user ID of the connecting process.

30 `PMIX_GRPID` "pmix.egid" (`uint32_t`)

31 Effective group ID of the connecting process.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Request a credential from the PMIx server library or the host environment. The credential is returned as a **pmix_byte_object_t** to support potential binary formats - it is therefore opaque to the caller. No information as to the source of the credential is provided.

16.1.2 PMIx_Get_credential_nb

Summary

Request a credential from the PMIx server library or the host environment.

Format

PMIx v3.0

C

pmix_status_t

```
PMIx_Get_credential_nb(const pmix_info_t info[], size_t ninfo,  
                      pmix_credential_cbfunc_t cbfunc,  
                      void *cbdata);
```

C

IN info

Array of **pmix_info_t** structures (array of handles)

IN ninfo

Number of elements in the *info* array (**size_t**)

IN cbfunc

Callback function to return credential (**pmix_credential_cbfunc_t** function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed successfully - the *cbfunc* will *not* be called.

1 If none of the above return codes are appropriate, then an implementation must return either a
2 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

3 There are no required attributes for this API. Note that implementations may choose to internally
4 execute integration for some security environments (e.g., directly contacting a *munge* server).

5 Implementations that support the operation but cannot directly process the client's request must
6 pass any attributes that are provided by the client to the host environment for processing. In
7 addition, the following attributes are required to be included in the *info* array passed from the PMIx
8 library to the host environment:

9 **PMIX_USERID** "pmix.euid" (uint32_t)

10 Effective user ID of the connecting process.

11 **PMIX_GRPID** "pmix.egid" (uint32_t)

12 Effective group ID of the connecting process.

Optional Attributes

13 The following attributes are optional for host environments that support this operation:

14 **PMIX_TIMEOUT** "pmix.timeout" (int)

15 Time in seconds before the specified operation should time out (zero indicating infinite) and
16 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
17 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

18 Request a credential from the PMIx server library or the host environment. This version of the API
19 is generally preferred in scenarios where the host environment may have to contact a remote
20 credential service. Thus, provision is made for the system to return additional information (e.g., the
21 identity of the issuing agent) outside of the credential itself and visible to the application.
22

16.1.3 Credential Attributes

24 The following attributes are defined to support credential operations:

25 **PMIX_CRED_TYPE** "pmix.sec.ctype" (char*)

26 When passed in **PMIx_Get_credential**, a prioritized, comma-delimited list of desired
27 credential types for use in environments where multiple authentication mechanisms may be
28 available. When returned in a callback function, a string identifier of the credential type.

29 **PMIX_CRYPTO_KEY** "pmix.sec.key" (pmix_byte_object_t)

30 Blob containing crypto key.

1 16.2 Validating Credentials

2 Given a credential, PMIx provides two methods by which a caller can request that the system
3 validate it, returning any additional information (e.g., authorizations) conveyed within the
4 credential.

5 16.2.1 PMIx_Validate_credential

6 Summary

7 Request validation of a credential by the PMIx server library or the host environment.

8 Format

PMIx v3.0

C

```
9 pmix_status_t  
10 PMIx_Validate_credential(const pmix_byte_object_t *cred,  
11                          const pmix_info_t info[], size_t ninfo,  
12                          pmix_info_t **results, size_t *nresults);
```

C

13 IN cred

14 Pointer to [pmix_byte_object_t](#) containing the credential (handle)

15 IN info

16 Array of [pmix_info_t](#) structures (array of handles)

17 IN ninfo

18 Number of elements in the *info* array ([size_t](#))

19 INOUT results

20 Address where a pointer to an array of [pmix_info_t](#) containing the results of the request
21 can be returned (memory reference)

22 INOUT nresults

23 Address where the number of elements in *results* can be returned (handle)

24 A successful return indicates that the credential was valid and any information it contained was
25 successfully processed. Details of the result will be returned in the *results* array.

26 Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Required Attributes

27 There are no required attributes for this API. Note that implementations may choose to internally
28 execute integration for some security environments (e.g., directly contacting a *munge* server).

29 Implementations that support the operation but cannot directly process the client's request must
30 pass any attributes that are provided by the client to the host environment for processing. In
31 addition, the following attributes are required to be included in the *info* array passed from the PMIx
32 library to the host environment:

33 [PMIX_USERID](#) "pmix.euid" ([uint32_t](#))

1 Effective user ID of the connecting process.

2 **PMIX_GRPID** "pmix.egid" (uint32_t)

3 Effective group ID of the connecting process.



▼----- Optional Attributes -----▼

4 The following attributes are optional for host environments that support this operation:

5 **PMIX_TIMEOUT** "pmix.timeout" (int)

6 Time in seconds before the specified operation should time out (zero indicating infinite) and
7 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
8 caused by multiple layers (client, server, and host) simultaneously timing the operation.



9 **Description**

10 Request validation of a credential by the PMIx server library or the host environment.

11 **16.2.2 PMIx_Validate_credential_nb**

12 **Summary**

13 Request validation of a credential by the PMIx server library or the host environment. Provision is
14 made for the system to return additional information regarding possible authorization limitations
15 beyond simple authentication.

Format

C

```
pmix_status_t
PMIx_Validate_credential_nb(const pmix_byte_object_t *cred,
                           const pmix_info_t info[], size_t ninfo,
                           pmix_validation_cbfunc_t cbfunc,
                           void *cbdata);
```

C

IN cred

Pointer to `pmix_byte_object_t` containing the credential (handle)

IN info

Array of `pmix_info_t` structures (array of handles)

IN ninfo

Number of elements in the *info* array (`size_t`)

IN cbfunc

Callback function to return result (`pmix_validation_cbfunc_t` function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when `PMIX_SUCCESS` is returned.

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

`PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed successfully - the *cbfunc* will *not* be called.

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

There are no required attributes for this API. Note that implementations may choose to internally execute integration for some security environments (e.g., directly contacting a *munge* server).

Implementations that support the operation but cannot directly process the client's request must pass any attributes that are provided by the client to the host environment for processing. In addition, the following attributes are required to be included in the *info* array passed from the PMIx library to the host environment:

`PMIX_USERID` "pmix.euid" (`uint32_t`)

Effective user ID of the connecting process.

`PMIX_GRPID` "pmix.egid" (`uint32_t`)

Effective group ID of the connecting process.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Request validation of a credential by the PMIx server library or the host environment. This version of the API is generally preferred in scenarios where the host environment may have to contact a remote credential service. Provision is made for the system to return additional information (e.g., possible authorization limitations) beyond simple authentication.

CHAPTER 17

Server-Specific Interfaces

1 The process that hosts the PMIx server library interacts with that library in two distinct manners.
2 First, PMIx provides a set of APIs by which the host can request specific services from its library.
3 This includes:

- 4 • collecting inventory to support scheduling algorithms,
- 5 • providing subsystems with an opportunity to precondition their resources for optimized
6 application support,
- 7 • generating regular expressions,
- 8 • registering information to be passed to client processes, and
- 9 • requesting information on behalf of a remote process.

10 Note that the host always has access to all PMIx client APIs - the functions listed below are in
11 addition to those available to a PMIx client.

12 Second, the host can provide a set of callback functions by which the PMIx server library can pass
13 requests upward for servicing by the host. These include notifications of client connection and
14 finalize, as well as requests by clients for information and/or services that the PMIx server library
15 does not itself provide.

17.1 Server Initialization and Finalization

17 Initialization and finalization routines for PMIx servers.

17.1.1 PMIx_server_init

Summary

Initialize the PMIx server.

Format

PMIx v1.0

C

```
pmix_status_t  
PMIx_server_init(pmix_server_module_t *module,  
                pmix_info_t info[], size_t ninfo);
```

INOUT module

pmix_server_module_t structure (handle)
IN info
 Array of **pmix_info_t** structures (array of handles)
IN ninfo
 Number of elements in the *info* array (**size_t**)

Returns **PMIX_SUCCESS** or a negative value indicating the error.

----- Required Attributes -----

The following attributes are required to be supported by all PMIx libraries:

PMIX_SERVER_NAMESPACE "pmix.srv.namespace" (**char***)

Name of the namespace to use for this PMIx server.

PMIX_SERVER_RANK "pmix.srv.rank" (**pmix_rank_t**)

Rank of this PMIx server.

PMIX_SERVER_TMPDIR "pmix.srvr.tmpdir" (**char***)

Top-level temporary directory for all client processes connected to this server, and where the PMIx server will place its tool rendezvous point and contact information.

PMIX_SYSTEM_TMPDIR "pmix.sys.tmpdir" (**char***)

Temporary directory for this system, and where a PMIx server that declares itself to be a system-level server will place a tool rendezvous point and contact information.

PMIX_SERVER_TOOL_SUPPORT "pmix.srvr.tool" (**bool**)

The host RM wants to declare itself as willing to accept tool connection requests.

PMIX_SERVER_SYSTEM_SUPPORT "pmix.srvr.sys" (**bool**)

The host RM wants to declare itself as being the local system server for PMIx connection requests.

PMIX_SERVER_SESSION_SUPPORT "pmix.srvr.sess" (**bool**)

The host RM wants to declare itself as being the local session server for PMIx connection requests.

PMIX_SERVER_GATEWAY "pmix.srv.gway" (**bool**)

Server is acting as a gateway for PMIx requests that cannot be serviced on backend nodes (e.g., logging to email).

PMIX_SERVER_SCHEDULER "pmix.srv.sched" (**bool**)

Server is supporting system scheduler and desires access to appropriate WLM-supporting features. Indicates that the library is to be initialized for scheduler support.

Optional Attributes

The following attributes are optional for implementers of PMIx libraries:

PMIX_USOCK_DISABLE "pmix.usock.disable" (bool)

Disable legacy UNIX socket (usock) support. If the library supports Unix socket connections, this attribute may be supported for disabling it.

PMIX_SOCKET_MODE "pmix.sockmode" (uint32_t)

POSIX *mode_t* (9 bits valid). If the library supports socket connections, this attribute may be supported for setting the socket mode.

PMIX_SINGLE_LISTENER "pmix.sing.listnr" (bool)

Use only one rendezvous socket, letting priorities and/or environment parameters select the active transport.

PMIX_TCP_REPORT_URI "pmix.tcp.repuri" (char*)

If provided, directs that the TCP URI be reported and indicates the desired method of reporting: '-' for stdout, '+' for stderr, or filename. If the library supports TCP socket connections, this attribute may be supported for reporting the URI.

PMIX_TCP_IF_INCLUDE "pmix.tcp.ifinclude" (char*)

Comma-delimited list of devices and/or CIDR notation to include when establishing the TCP connection. If the library supports TCP socket connections, this attribute may be supported for specifying the interfaces to be used.

PMIX_TCP_IF_EXCLUDE "pmix.tcp.ifexclude" (char*)

Comma-delimited list of devices and/or CIDR notation to exclude when establishing the TCP connection. If the library supports TCP socket connections, this attribute may be supported for specifying the interfaces that are *not* to be used.

PMIX_TCP_IPV4_PORT "pmix.tcp.ipv4" (int)

The IPv4 port to be used.. If the library supports IPV4 connections, this attribute may be supported for specifying the port to be used.

PMIX_TCP_IPV6_PORT "pmix.tcp.ipv6" (int)

The IPv6 port to be used. If the library supports IPV6 connections, this attribute may be supported for specifying the port to be used.

PMIX_TCP_DISABLE_IPV4 "pmix.tcp.disipv4" (bool)

Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections, this attribute may be supported for disabling it.

PMIX_TCP_DISABLE_IPV6 "pmix.tcp.disipv6" (bool)

Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections, this attribute may be supported for disabling it.

PMIX_SERVER_REMOTE_CONNECTIONS "pmix.srvr.remote" (bool)

1 Allow connections from remote tools. Forces the PMIx server to not exclusively use
2 loopback device. If the library supports connections from remote tools, this attribute may
3 be supported for enabling or disabling it.

4 **PMIX_EXTERNAL_PROGRESS** "pmix.evevt" (bool)

5 The host shall progress the PMIx library via calls to **PMIx_Progress**

6 **PMIX_EVENT_BASE** "pmix.evbase" (void*)

7 Pointer to an **event_base** to use in place of the internal progress thread. All PMIx library
8 events are to be assigned to the provided event base. The event base *must* be compatible with
9 the event library used by the PMIx implementation - e.g., either both the host and PMIx
10 library must use libevent, or both must use libev. Cross-matches are unlikely to work and
11 should be avoided - it is the responsibility of the host to ensure that the PMIx
12 implementation supports (and was built with) the appropriate event library.

13 **PMIX_TOPOLOGY2** "pmix.topo2" (pmix_topology_t)

14 Provide a pointer to an implementation-specific description of the local node topology.

15 **PMIX_SERVER_SHARE_TOPOLOGY** "pmix.srvr.share" (bool)

16 The PMIx server is to share its copy of the local node topology (whether given to it or
17 self-discovered) with any clients. The PMIx server will perform the necessary actions to
18 scalably expose the description to the local clients. This includes creating any required
19 shared memory backing stores and/ or XML representations, plus ensuring that all necessary
20 key-value pairs for clients to access the description are included in the job-level information
21 provided to each client. All required files are to be installed under the effective
22 **PMIX_SERVER_TMPDIR** directory. The PMIx server library is responsible for cleaning up
23 any artifacts (e.g., shared memory backing files or cached key-value pairs) at library finalize.

24 **PMIX_SERVER_ENABLE_MONITORING** "pmix.srv.monitor" (bool)

25 Enable PMIx internal monitoring by the PMIx server.

26 **PMIX_HOMOGENEOUS_SYSTEM** "pmix.homo" (bool)

27 The nodes comprising the session are homogeneous - i.e., they each contain the same
28 number of identical packages, fabric interfaces, GPUs, and other devices.

29 **PMIX_SINGLETON** "pmix.singleton" (char*)

30 String representation (nspace.rank) of proc ID for the singleton the server was started to
31 support

32 **PMIX_IOF_LOCAL_OUTPUT** "pmix.iof.local" (bool)

33 Write output streams to local stdout/err

▲-----▲

Description

Initialize the PMIx server support library, and provide a pointer to a `pmix_server_module_t` structure containing the caller's callback functions. The array of `pmix_info_t` structs is used to pass additional info that may be required by the server when initializing. For example, it may include the `PMIX_SERVER_TOOL_SUPPORT` attribute, thereby indicating that the daemon is willing to accept connection requests from tools.

Advice to PMIx server hosts

Providing a value of `NULL` for the *module* argument is permitted, as is passing an empty *module* structure. Doing so indicates that the host environment will not provide support for multi-node operations such as `PMIx_Fence`, but does intend to support local clients access to information.

17.1.2 `PMIx_server_finalize`

Summary

Finalize the PMIx server library.

Format

PMIx v1.0

```
pmix_status_t  
PMIx_server_finalize(void);
```

Returns `PMIX_SUCCESS` or a negative value indicating the error.

Description

Finalize the PMIx server support library, terminating all connections to attached tools and any local clients. All memory usage is released.

1 17.1.3 Server Initialization Attributes

2 These attributes are used to direct the configuration and operation of the PMIx server library by
3 passing them into `PMIx_server_init`.

4 **PMIX_TOPOLOGY2** "`pmix.topo2`" (`pmix_topology_t`)

5 Provide a pointer to an implementation-specific description of the local node topology.

6 **PMIX_SERVER_SHARE_TOPOLOGY** "`pmix.srvr.share`" (`bool`)

7 The PMIx server is to share its copy of the local node topology (whether given to it or
8 self-discovered) with any clients.

9 **PMIX_USOCK_DISABLE** "`pmix.usock.disable`" (`bool`)

10 Disable legacy UNIX socket (usock) support.

11 **PMIX_SOCKET_MODE** "`pmix.sockmode`" (`uint32_t`)

12 POSIX `mode_t` (9 bits valid).

13 **PMIX_SINGLE_LISTENER** "`pmix.sing.listnr`" (`bool`)

14 Use only one rendezvous socket, letting priorities and/or environment parameters select the
15 active transport.

16 **PMIX_SERVER_TOOL_SUPPORT** "`pmix.srvr.tool`" (`bool`)

17 The host RM wants to declare itself as willing to accept tool connection requests.

18 **PMIX_SERVER_REMOTE_CONNECTIONS** "`pmix.srvr.remote`" (`bool`)

19 Allow connections from remote tools. Forces the PMIx server to not exclusively use
20 loopback device.

21 **PMIX_SERVER_SYSTEM_SUPPORT** "`pmix.srvr.sys`" (`bool`)

22 The host RM wants to declare itself as being the local system server for PMIx connection
23 requests.

24 **PMIX_SERVER_SESSION_SUPPORT** "`pmix.srvr.sess`" (`bool`)

25 The host RM wants to declare itself as being the local session server for PMIx connection
26 requests.

27 **PMIX_SERVER_START_TIME** "`pmix.srvr.strtime`" (`char*`)

28 Time when the server started - i.e., when the server created its rendezvous file (given in
29 ctime string format).

30 **PMIX_SERVER_TMPDIR** "`pmix.srvr.tmpdir`" (`char*`)

31 Top-level temporary directory for all client processes connected to this server, and where the
32 PMIx server will place its tool rendezvous point and contact information.

33 **PMIX_SYSTEM_TMPDIR** "`pmix.sys.tmpdir`" (`char*`)

34 Temporary directory for this system, and where a PMIx server that declares itself to be a
35 system-level server will place a tool rendezvous point and contact information.

36 **PMIX_SERVER_ENABLE_MONITORING** "`pmix.srv.monitor`" (`bool`)

37 Enable PMIx internal monitoring by the PMIx server.

38 **PMIX_SERVER_NAMESPACE** "`pmix.srv.namespace`" (`char*`)

39 Name of the namespace to use for this PMIx server.

40 **PMIX_SERVER_RANK** "`pmix.srv.rank`" (`pmix_rank_t`)

41 Rank of this PMIx server.

42 **PMIX_SERVER_GATEWAY** "`pmix.srv.gway`" (`bool`)

1 Server is acting as a gateway for PMIx requests that cannot be serviced on backend nodes
2 (e.g., logging to email).

3 **PMIX_SERVER_SCHEDULER** "pmix.srv.sched" (bool)

4 Server is supporting system scheduler and desires access to appropriate WLM-supporting
5 features. Indicates that the library is to be initialized for scheduler support.

6 **PMIX_EXTERNAL_PROGRESS** "pmix.evext" (bool)

7 The host shall progress the PMIx library via calls to **PMIx_Progress**

8 **PMIX_HOMOGENEOUS_SYSTEM** "pmix.homo" (bool)

9 The nodes comprising the session are homogeneous - i.e., they each contain the same
10 number of identical packages, fabric interfaces, GPUs, and other devices.

11 *Provisional* **PMIX_SINGLETON** "pmix.singleton" (char*)

12 String representation (nspace.rank) of proc ID for the singleton the server was started to
13 support

14 17.2 Server Support Functions

15 The following APIs allow the RM daemon that hosts the PMIx server library to request specific
16 services from the PMIx library.

17 17.2.1 PMIx_generate_regex

18 Summary

19 Generate a compressed representation of the input string.

20 Format

PMIx v1.0

21 pmix_status_t

22 PMIx_generate_regex(const char *input, char **output);

23 **IN** input

String to process (string)

25 **OUT** output

26 Compressed representation of *input* (array of bytes)

27 Returns **PMIX_SUCCESS** or a negative value indicating the error.

Description

Given a comma-separated list of *input* values, generate a reduced size representation of the input that can be passed down to the PMIx server library's `PMIx_server_register_nspace` API for parsing. The order of the individual values in the *input* string is preserved across the operation. The caller is responsible for releasing the returned data.

The precise compressed representations will be implementation specific. The regular expression itself is not required to be a printable string nor to obey typical string constraints (e.g., include a **NULL** terminator byte). However, all PMIx implementations are required to include a colon-delimited **NULL**-terminated string at the beginning of the output representation that can be printed for diagnostic purposes and identifies the method used to generate the representation. The following identifiers are reserved by the PMIx Standard:

- "**raw**:\0" - indicates that the expression following the identifier is simply the comma-delimited input string (no processing was performed).
- "**pmix**:\0" - a PMIx-unique regular expression represented as a **NULL**-terminated string following the identifier.
- "**blob**:\0" - a PMIx-unique regular expression that is not represented as a **NULL**-terminated string following the identifier. Additional implementation-specific metadata may follow the identifier along with the data itself. For example, a compressed binary array format based on the *zlib* compression package, with the size encoded in the space immediately following the identifier.

Communicating the resulting output should be done by first packing the returned expression using the `PMIx_Data_pack`, declaring the input to be of type `PMIX_REGEX`, and then obtaining the resulting blob to be communicated using the `PMIX_DATA_BUFFER_UNLOAD` macro. The reciprocal method can be used on the remote end prior to passing the regex into `PMIx_server_register_nspace`. The pack/unpack routines will ensure proper handling of the data based on the regex prefix.

17.2.2 PMIx_generate_ppn

Summary

Generate a compressed representation of the input identifying the processes on each node.

Format

PMIx v1.0

```
pmix_status_t  
PMIx_generate_ppn(const char *input, char **ppn);
```

IN `input`
String to process (string)

1 **OUT** `ppn`
2 Compressed representation of *input* (array of bytes)
3 Returns **PMIX_SUCCESS** or a negative value indicating the error.

4 **Description**

5 The input shall consist of a semicolon-separated list of ranges representing the ranks of processes
6 on each node of the job - e.g., "1-4;2-5;8,10,11,12;6,7,9". Each field of the input must
7 correspond to the node name provided at that position in the input to **PMIX_generate_regex**.
8 Thus, in the example, ranks 1-4 would be located on the first node of the comma-separated list of
9 names provided to **PMIX_generate_regex**, and ranks 2-5 would be on the second name in the
10 list.

11 Rules governing the format of the returned regular expression are the same as those specified for
12 **PMIX_generate_regex**, as detailed [here](#).

13 **17.2.3 PMIX_server_register_namespace**

14 **Summary**

15 Setup the data about a particular namespace.

16 **Format**

PMIX v1.0

C

```
17 pmix_status_t  
18 PMIX_server_register_namespace(const pmix_namespace_t nspace,  
19                               int nlocalprocs,  
20                               pmix_info_t info[], size_t ninfo,  
21                               pmix_op_cbfunc_t cbfunc,  
22                               void *cbdata);
```

C

23 **IN** **nspace**
24 Character array of maximum size **PMIX_MAX_NSLEN** containing the namespace identifier
25 (string)
26 **IN** **nlocalprocs**
27 number of local processes (integer)
28 **IN** **info**
29 Array of info structures (array of handles)
30 **IN** **ninfo**
31 Number of elements in the *info* array (integer)
32 **IN** **cbfunc**
33 Callback function **pmix_op_cbfunc_t** to be executed upon completion of the operation.
34 A **NULL** function reference indicates that the function is to be executed as a blocking
35 operation (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning from the API. The callback function, `cbfunc`, is only called when `PMIX_SUCCESS` is returned.

Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- `PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed and returned *success* - the `cbfunc` will not be called

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

`PMIX_REGISTER_NODATA` "pmix.reg.nodata" (`bool`)

Registration is for this namespace only, do not copy job data.

`PMIX_SESSION_INFO_ARRAY` "pmix.ssn.arr" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing session-realm information. The `PMIX_SESSION_ID` attribute is required to be included in the array.

`PMIX_JOB_INFO_ARRAY` "pmix.job.arr" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing job-realm information. The `PMIX_SESSION_ID` attribute of the *session* containing the *job* is required to be included in the array whenever the PMIx server library may host multiple sessions (e.g., when executing with a host RM daemon). As information is registered one job (aka namespace) at a time via the `PMIx_server_register_namespace` API, there is no requirement that the array contain either the `PMIX_NAMESPACE` or `PMIX_JOBID` attributes when used in that context (though either or both of them may be included). At least one of the job identifiers must be provided in all other contexts where the job being referenced is ambiguous.

`PMIX_APP_INFO_ARRAY` "pmix.app.arr" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing application-realm information. The `PMIX_NAMESPACE` or `PMIX_JOBID` attributes of the *job* containing the application, plus its `PMIX_APPNUM` attribute, must be included in the array when the array is *not* included as part of a call to `PMIx_server_register_namespace` - i.e., when the job containing the application is ambiguous. The job identification is otherwise optional.

`PMIX_PROC_INFO_ARRAY` "pmix.pdata" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing process-realm information. The `PMIX_RANK` and `PMIX_NAMESPACE` attributes, or the `PMIX_PROCID` attribute, are required to be included in the array when the array is not included as part of a call to `PMIx_server_register_namespace` - i.e., when the job containing the process is ambiguous. All three may be included if desired. When the array is included in some

broader structure that identifies the job, then only the `PMIX_RANK` or the `PMIX_PROCID` attribute must be included (the others are optional).

`PMIX_NODE_INFO_ARRAY` "pmix.node.arr" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing node-realm information. At a minimum, either the `PMIX_NODEID` or `PMIX_HOSTNAME` attribute is required to be included in the array, though both may be included.

Host environments are required to provide a wide range of session-, job-, application-, node-, and process-realm information, and may choose to provide a similarly wide range of optional information. The information is broadly separated into categories based on the *data realm* definitions explained in Section 6.1, and retrieved according to the rules detailed in Section 6.2.

Session-realm information may be passed as individual `pmix_info_t` entries, or as part of a `pmix_data_array_t` using the `PMIX_SESSION_INFO_ARRAY` attribute. The list of data referenced in this way shall include:

- `PMIX_UNIV_SIZE` "pmix.univ.size" (`uint32_t`)
Maximum number of process that can be simultaneously executing in a session. Note that this attribute is equivalent to the `PMIX_MAX_PROCS` attribute for the *session* realm - it is included in the PMIx Standard for historical reasons.
- `PMIX_MAX_PROCS` "pmix.max.size" (`uint32_t`)
Maximum number of processes that can be simultaneously executed in the specified realm. Typically, this is a constraint imposed by a scheduler or by user settings in a hostfile or other resource description. Defaults to the *job* realm. Must be provided if `PMIX_UNIV_SIZE` is not given. Requires use of the `PMIX_SESSION_INFO` attribute to avoid ambiguity when retrieving it.
- `PMIX_SESSION_ID` "pmix.session.id" (`uint32_t`)
Session identifier assigned by the scheduler.

plus the following optional information:

- `PMIX_CLUSTER_ID` "pmix.clid" (`char*`)
A string name for the cluster this allocation is on. As this information is not related to the namespace, it is best passed using the `PMIx_server_register_resources` API.
- `PMIX_ALLOCATED_NODELIST` "pmix.alist" (`char*`)
Comma-delimited list or regular expression of all nodes in the specified realm regardless of whether or not they currently host processes. Defaults to the *job* realm.
- `PMIX_RM_NAME` "pmix.rm.name" (`char*`)
String name of the RM. As this information is not related to the namespace, it is best passed using the `PMIx_server_register_resources` API.
- `PMIX_RM_VERSION` "pmix.rm.version" (`char*`)

1 RM version string. As this information is not related to the namespace, it is best passed
2 using the [PMIx_server_register_resources](#) API.

- 3 • [PMIX_SERVER_HOSTNAME](#) "pmix.srvr.host" (char*)
4 Host where target PMIx server is located. As this information is not related to the
5 namespace, it is best passed using the [PMIx_server_register_resources](#) API.

6 Job-realm information may be passed as individual [pmix_info_t](#) entries, or as part of a
7 [pmix_data_array_t](#) using the [PMIX_JOB_INFO_ARRAY](#) attribute. The list of data
8 referenced in this way shall include:

- 9 • [PMIX_SERVER_NAMESPACE](#) "pmix.srv.namespace" (char*)
10 Name of the namespace to use for this PMIx server. Identifies the namespace of the PMIx
11 server itself
- 12 • [PMIX_SERVER_RANK](#) "pmix.srv.rank" (pmix_rank_t)
13 Rank of this PMIx server. Identifies the rank of the PMIx server itself.
- 14 • [PMIX_NAMESPACE](#) "pmix.namespace" (char*)
15 Namespace of the job - may be a numerical value expressed as a string, but is often an
16 alphanumeric string carrying information solely of use to the system. Required to be
17 unique within the scope of the host environment. One cannot retrieve the namespace of an
18 arbitrary process since that would require already knowing the namespace of that process.
19 However, a process' own namespace can be retrieved by passing a NULL value of *proc* to
20 [PMIx_Get](#). Identifies the namespace of the job being registered.
- 21 • [PMIX_JOBID](#) "pmix.jobid" (char*)
22 Job identifier assigned by the scheduler to the specified job - may be identical to the
23 namespace, but is often a numerical value expressed as a string (e.g., "12345.3").
- 24 • [PMIX_JOB_SIZE](#) "pmix.job.size" (uint32_t)
25 Total number of processes in the specified job across all contained applications. Note that
26 this value can be different from [PMIX_MAX_PROCS](#). For example, users may choose to
27 subdivide an allocation (running several jobs in parallel within it), and dynamic
28 programming models may support adding and removing processes from a running *job*
29 on-the-fly. In the latter case, PMIx events may be used to notify processes within the job
30 that the job size has changed.
- 31 • [PMIX_MAX_PROCS](#) "pmix.max.size" (uint32_t)
32 Maximum number of processes that can be simultaneously executed in the specified realm.
33 Typically, this is a constraint imposed by a scheduler or by user settings in a hostfile or
34 other resource description. Defaults to the *job* realm. Retrieval of this attribute defaults to
35 the job level unless an appropriate specification is given (e.g., [PMIX_SESSION_INFO](#)).
- 36 • [PMIX_NODE_MAP](#) "pmix.nmap" (char*)
37 Regular expression of nodes currently hosting processes in the specified realm - see
38 [17.2.3.2](#) for an explanation of its generation. Defaults to the *job* realm.

- 1 • **PMIX_PROC_MAP** "pmix.pmap" (char*)
2 Regular expression describing processes on each node in the specified realm - see 17.2.3.2
3 for an explanation of its generation. Defaults to the *job* realm.

4 plus the following optional information:

- 5 • **PMIX_NPROC_OFFSET** "pmix.offset" (pmix_rank_t)
6 Starting global rank of the specified job. The returned value is the same as the value of
7 **PMIX_GLOBAL_RANK** of rank 0 of the specified job.
- 8 • **PMIX_JOB_NUM_APPS** "pmix.job.napps" (uint32_t)
9 Number of applications in the specified job. This is a required attribute if more than one
10 application is included in the job.
- 11 • **PMIX_MAPBY** "pmix.mapby" (char*)
12 Process mapping policy - when accessed using **PMIx_Get**, use the
13 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
14 provided namespace. Supported values are launcher specific.
- 15 • **PMIX_RANKBY** "pmix.rankby" (char*)
16 Process ranking policy - when accessed using **PMIx_Get**, use the
17 **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for
18 the provided namespace. Supported values are launcher specific.
- 19 • **PMIX_BINDTO** "pmix.bindto" (char*)
20 Process binding policy - when accessed using **PMIx_Get**, use the
21 **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the
22 provided namespace. Supported values are launcher specific.
- 23 • **PMIX_HOSTNAME_KEEP_FQDN** "pmix.fqdn" (bool)
24 FQDNs are being retained by the PMIx library.
- 25 • **PMIX_ANL_MAP** "pmix.anlmap" (char*)
26 Process map equivalent to **PMIX_PROC_MAP** expressed in Argonne National
27 Laboratory's PMI-1/PMI-2 notation. Defaults to the *job* realm.
- 28 • **PMIX_TDIR_RMCLEAN** "pmix.tdir.rmclean" (bool)
29 The Resource Manager will remove any directories or files it creates in **PMIX_TMPDIR**.
- 30 • **PMIX_CRYPTO_KEY** "pmix.sec.key" (pmix_byte_object_t)
31 Blob containing crypto key.

32 If more than one application is included in the namespace, then the host environment is also
33 required to supply data consisting of the following items for each application in the job, passed as a
34 **pmix_data_array_t** using the **PMIX_APP_INFO_ARRAY** attribute:

- 35 • **PMIX_APPNUM** "pmix.appnum" (uint32_t)
36 The application number within the job in which the specified process is a member. This
37 attribute must appear at the beginning of the array.

- 1 • **PMIX_APP_SIZE** "pmix.app.size" (uint32_t)
2 Number of processes in the specified application, regardless of their execution state - i.e.,
3 this number may include processes that either failed to start or have already terminated.
- 4 • **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)
5 Maximum number of processes that can be simultaneously executed in the specified
6 realm. Typically, this is a constraint imposed by a scheduler or by user settings in a
7 hostfile or other resource description. Defaults to the *job* realm. Requires use of the
8 **PMIX_APP_INFO** attribute to avoid ambiguity when retrieving it.
- 9 • **PMIX_APPLDR** "pmix.aldr" (pmix_rank_t)
10 Lowest rank in the specified application.
- 11 • **PMIX_WDIR** "pmix.wdir" (char*)
12 Working directory for spawned processes. This attribute is required for all registrations,
13 but may be provided as an individual **pmix_info_t** entry if only one application is
14 included in the namespace.
- 15 • **PMIX_APP_ARGV** "pmix.app.argv" (char*)
16 Consolidated argv passed to the spawn command for the given application (e.g., "/myapp
17 arg1 arg2 arg3"). This attribute is required for all registrations, but may be provided as an
18 individual **pmix_info_t** entry if only one application is included in the namespace.

19 plus the following optional information:

- 20 • **PMIX_PSET_NAMES** "pmix.pset.nms" (pmix_data_array_t*)
21 Returns an array of **char*** string names of the process sets in which the given process is
22 a member.
- 23 • **PMIX_APP_MAP_TYPE** "pmix.apmap.type" (char*)
24 Type of mapping used to layout the application (e.g., **cyclic**). This attribute may be
25 provided as an individual **pmix_info_t** entry if only one application is included in the
26 namespace.
- 27 • **PMIX_APP_MAP_REGEX** "pmix.apmap.regex" (char*)
28 Regular expression describing the result of the process mapping. This attribute may be
29 provided as an individual **pmix_info_t** entry if only one application is included in the
30 namespace.

31 The data may also include attributes provided by the host environment that identify the
32 programming model (as specified by the user) being executed within the application. The PMIx
33 server library may utilize this information to customize the environment to fit that model (e.g.,
34 adding environmental variables specified by the corresponding standard for that model):

- 35 • **PMIX_PROGRAMMING_MODEL** "pmix.pgm.model" (char*)
36 Programming model being initialized (e.g., "MPI" or "OpenMP").
- 37 • **PMIX_MODEL_LIBRARY_NAME** "pmix.mdl.name" (char*)
38 Programming model implementation ID (e.g., "OpenMPI" or "MPICH").

- 1 • **PMIX_MODEL_LIBRARY_VERSION** "pmix.mld.vrs" (char*)

2 Programming model version string (e.g., "2.1.1").

3 Node-realm information may be passed as individual **pmix_info_t** entries if only one node will
4 host processes from the job being registered, or as part of a **pmix_data_array_t** using the
5 **PMIX_NODE_INFO_ARRAY** attribute when multiple nodes are involved in the job. The list of data
6 referenced in this way shall include:

- 7 • **PMIX_NODEID** "pmix.nodeid" (uint32_t)

8 Node identifier expressed as the node's index (beginning at zero) in an array of nodes
9 within the active session. The value must be unique and directly correlate to the
10 **PMIX_HOSTNAME** of the node - i.e., users can interchangeably reference the same
11 location using either the **PMIX_HOSTNAME** or corresponding **PMIX_NODEID**.

- 12 • **PMIX_HOSTNAME** "pmix.hname" (char*)

13 Name of the host, as returned by the **gethostname** utility or its equivalent. As this
14 information is not related to the namespace, it can be passed using the
15 **PMIx_server_register_resources** API. However, either it or the
16 **PMIX_NODEID** must be included in the array to properly identify the node.

- 17 • **PMIX_HOSTNAME_ALIASES** "pmix.alias" (char*)

18 Comma-delimited list of names by which the target node is known. As this information is
19 not related to the namespace, it is best passed using the
20 **PMIx_server_register_resources** API.

- 21 • **PMIX_LOCAL_SIZE** "pmix.local.size" (uint32_t)

22 Number of processes in the specified job or application on the caller's node. Defaults to
23 job unless the **PMIX_APP_INFO** and the **PMIX_APPNUM** qualifiers are given.

- 24 • **PMIX_NODE_SIZE** "pmix.node.size" (uint32_t)

25 Number of processes across all jobs that are executing upon the node.

- 26 • **PMIX_LOCALLDR** "pmix.lldr" (pmix_rank_t)

27 Lowest rank within the specified job on the node (defaults to current node in absence of
28 **PMIX_HOSTNAME** or **PMIX_NODEID** qualifier).

- 29 • **PMIX_LOCAL_PEERS** "pmix.lpeers" (char*)

30 Comma-delimited list of ranks that are executing on the local node within the specified
31 namespace – shortcut for **PMIx_Resolve_peers** for the local node.

- 32 • **PMIX_NODE_OVERSUBSCRIBED** "pmix.ndosub" (bool)

33 True if the number of processes from this job on this node exceeds the number of slots
34 allocated to it

35 plus the following information for the server's own node:

- 36 • **PMIX_TMPDIR** "pmix.tmpdir" (char*)

37 Full path to the top-level temporary directory assigned to the session.

- 1 • **PMIX_NSDIR** "pmix.nmdir" (char*)
- 2 Full path to the temporary directory assigned to the specified job, under **PMIX_TMPDIR**.
- 3 • **PMIX_LOCAL_PROCS** "pmix.lprocs" (pmix_proc_t array)
- 4 Array of **pmix_proc_t** of all processes executing on the local node – shortcut for
- 5 **PMIx_Resolve_peers** for the local node and a **NULL** namespace argument. The
- 6 process identifier is ignored for this attribute.

7 The data may also include the following optional information for the server's own node:

- 8 • **PMIX_LOCAL_CPUSETS** "pmix.lcpus" (pmix_data_array_t)
- 9 A **pmix_data_array_t** array of string representations of the PU binding bitmaps
- 10 applied to each local *peer* on the caller's node upon launch. Each string shall begin with
- 11 the name of the library that generated it (e.g., "hwloc") followed by a colon and the bitmap
- 12 string itself. The array shall be in the same order as the processes returned by
- 13 **PMIX_LOCAL_PEERS** for that namespace.
- 14 • **PMIX_AVAIL_PHYS_MEMORY** "pmix.pmem" (uint64_t)
- 15 Total available physical memory on a node. As this information is not related to the
- 16 namespace, it can be passed using the **PMIx_server_register_resources** API.

17 and the following optional information for other nodes:

- 18 • **PMIX_MAX_PROCS** "pmix.max.size" (uint32_t)
- 19 Maximum number of processes that can be simultaneously executed in the specified
- 20 realm. Typically, this is a constraint imposed by a scheduler or by user settings in a
- 21 hostfile or other resource description. Defaults to the *job* realm. Requires use of the
- 22 **PMIX_NODE_INFO** attribute to avoid ambiguity when retrieving it.

23 Process-realm information shall include the following data for each process in the job, passed as a

24 **pmix_data_array_t** using the **PMIX_PROC_INFO_ARRAY** attribute:

- 25 • **PMIX_RANK** "pmix.rank" (pmix_rank_t)
- 26 Process rank within the job, starting from zero.
- 27 • **PMIX_APPNUM** "pmix.appnum" (uint32_t)
- 28 The application number within the job in which the specified process is a member. This
- 29 attribute may be omitted if only one application is present in the namespace.
- 30 • **PMIX_APP_RANK** "pmix.apprank" (pmix_rank_t)
- 31 Rank of the specified process within its application. This attribute may be omitted if only
- 32 one application is present in the namespace.
- 33 • **PMIX_GLOBAL_RANK** "pmix.grank" (pmix_rank_t)
- 34 Rank of the specified process spanning across all jobs in this session, starting with zero.
- 35 Note that no ordering of the jobs is implied when computing this value. As jobs can start
- 36 and end at random times, this is defined as a continually growing number - i.e., it is not
- 37 dynamically adjusted as individual jobs and processes are started or terminated.

- 1 • **PMIX_LOCAL_RANK** "pmix.lrank" (uint16_t)
2 Rank of the specified process on its node - refers to the numerical location (starting from
3 zero) of the process on its node when counting only those processes from the same job
4 that share the node, ordered by their overall rank within that job.
- 5 • **PMIX_NODE_RANK** "pmix.nrank" (uint16_t)
6 Rank of the specified process on its node spanning all jobs- refers to the numerical location
7 (starting from zero) of the process on its node when counting all processes (regardless of
8 job) that share the node, ordered by their overall rank within the job. The value represents
9 a snapshot in time when the specified process was started on its node and is not
10 dynamically adjusted as processes from other jobs are started or terminated on the node.
- 11 • **PMIX_NODEID** "pmix.nodeid" (uint32_t)
12 Node identifier expressed as the node's index (beginning at zero) in an array of nodes
13 within the active session. The value must be unique and directly correlate to the
14 **PMIX_HOSTNAME** of the node - i.e., users can interchangeably reference the same
15 location using either the **PMIX_HOSTNAME** or corresponding **PMIX_NODEID**.
- 16 • **PMIX_REINCARNATION** "pmix.reinc" (uint32_t)
17 Number of times this process has been re-instantiated - i.e., a value of zero indicates that
18 the process has never been restarted.
- 19 • **PMIX_SPAWNED** "pmix.spawned" (bool)
20 **true** if this process resulted from a call to **PMIx_Spawn**. Lack of inclusion (i.e., a return
21 status of **PMIX_ERR_NOT_FOUND**) corresponds to a value of **false** for this attribute.

22 plus the following information for processes that are local to the server:

- 23 • **PMIX_LOCALITY_STRING** "pmix.locstr" (char*)
24 String describing a process's bound location - referenced using the process's rank. The
25 string is prefixed by the implementation that created it (e.g., "hwloc") followed by a colon.
26 The remainder of the string represents the corresponding locality as expressed by the
27 underlying implementation. The entire string must be passed to
28 **PMIx_Get_relative_locality** for processing. Note that hosts are only required to
29 provide locality strings for local client processes - thus, a call to **PMIx_Get** for the
30 locality string of a process that returns **PMIX_ERR_NOT_FOUND** indicates that the
31 process is not executing on the same node.
- 32 • **PMIX_PROCDIR** "pmix.pdir" (char*)
33 Full path to the subdirectory under **PMIX_NSDIR** assigned to the specified process.
- 34 • **PMIX_PACKAGE_RANK** "pmix.pkgrank" (uint16_t)
35 Rank of the specified process on the *package* where this process resides - refers to the
36 numerical location (starting from zero) of the process on its package when counting only
37 those processes from the same job that share the package, ordered by their overall rank
38 within that job. Note that processes that are not bound to PUs within a single specific
39 package cannot have a package rank.

1 and the following optional information - note that some of this information can be derived from
2 information already provided by other attributes, but it may be included here for ease of retrieval by
3 users:

- 4 • **PMIX_HOSTNAME** "pmix.hname" (char*)
5 Name of the host, as returned by the `gethostname` utility or its equivalent.
- 6 • **PMIX_CPUSSET** "pmix.cpuset" (char*)
7 A string representation of the PU binding bitmap applied to the process upon launch. The
8 string shall begin with the name of the library that generated it (e.g., "hwloc") followed by
9 a colon and the bitmap string itself.
- 10 • **PMIX_CPUSSET_BITMAP** "pmix.bitmap" (pmix_cpuset_t*)
11 Bitmap applied to the process upon launch.
- 12 • **PMIX_DEVICE_DISTANCES** "pmix.dev.dist" (pmix_data_array_t)
13 Return an array of `pmix_device_distance_t` containing the minimum and
14 maximum distances of the given process location to all devices of the specified type on the
15 local node.

16
17 Attributes not directly provided by the host environment may be derived by the PMIx server library
18 from other required information and included in the data made available to the server library's
19 clients.

20 Description

21 Pass job-related information to the PMIx server library for distribution to local client processes.

Advice to PMIx server hosts

22 Host environments are required to execute this operation prior to starting any local application
23 process within the given namespace.

24 The PMIx server must register all namespaces that will participate in collective operations with
25 local processes. This means that the server must register a namespace even if it will not host any
26 local processes from within that namespace if any local process of another namespace might at
27 some point perform an operation involving one or more processes from the new namespace. This is
28 necessary so that the collective operation can identify the participants and know when it is locally
29 complete.

30 The caller must also provide the number of local processes that will be launched within this
31 namespace. This is required for the PMIx server library to correctly handle collectives as a
32 collective operation call can occur before all the local processes have been started.

33 A **NULL** *cbfunc* reference indicates that the function is to be executed as a blocking operation.

The number of local processes for any given namespace is generally fixed at the time of application launch. Calls to `PMIx_Spawn` result in processes launched in their own namespace, not that of their parent. However, it is possible for processes to *migrate* to another node via a call to `PMIx_Job_control_nb`, thus resulting in a change to the number of local processes on both the initial node and the node to which the process moved. It is therefore critical that applications not migrate processes without first ensuring that PMIx-based collective operations are not in progress, and that no such operations be initiated until process migration has completed.

8 17.2.3.1 Namespace registration attributes

The following attributes are defined specifically for use with the `PMIx_server_register_nspace` API: `PMIX_REGISTER_NODATA`
"pmix.reg.nodata" (bool)
 Registration is for this namespace only, do not copy job data.

The following attributes are used to assemble information according to its data realm (*session*, *job*, *application*, *node*, or *process* as defined in Section 6.1) for registration where ambiguity may exist - see 17.2.3.2 for examples of their use.

PMIX_SESSION_INFO_ARRAY "pmix.ssn.arr" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing session-realm information. The `PMIX_SESSION_ID` attribute is required to be included in the array.

PMIX_JOB_INFO_ARRAY "pmix.job.arr" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing job-realm information. The `PMIX_SESSION_ID` attribute of the *session* containing the *job* is required to be included in the array whenever the PMIx server library may host multiple sessions (e.g., when executing with a host RM daemon). As information is registered one job (aka namespace) at a time via the `PMIx_server_register_nspace` API, there is no requirement that the array contain either the `PMIX_NAMESPACE` or `PMIX_JOBID` attributes when used in that context (though either or both of them may be included). At least one of the job identifiers must be provided in all other contexts where the job being referenced is ambiguous.

PMIX_APP_INFO_ARRAY "pmix.app.arr" (`pmix_data_array_t`)

Provide an array of `pmix_info_t` containing application-realm information. The `PMIX_NAMESPACE` or `PMIX_JOBID` attributes of the *job* containing the application, plus its `PMIX_APPNUM` attribute, must to be included in the array when the array is *not* included as part of a call to `PMIx_server_register_nspace` - i.e., when the job containing the application is ambiguous. The job identification is otherwise optional.

PMIX_PROC_INFO_ARRAY "pmix.pdata" (`pmix_data_array_t`)

1 Provide an array of `pmix_info_t` containing process-realm information. The
2 `PMIX_RANK` and `PMIX_NAMESPACE` attributes, or the `PMIX_PROCID` attribute, are required
3 to be included in the array when the array is not included as part of a call to
4 `PMIx_server_register_namespace` - i.e., when the job containing the process is
5 ambiguous. All three may be included if desired. When the array is included in some
6 broader structure that identifies the job, then only the `PMIX_RANK` or the `PMIX_PROCID`
7 attribute must be included (the others are optional).

8 **PMIX_NODE_INFO_ARRAY** "`pmix.node.arr`" (`pmix_data_array_t`)

9 Provide an array of `pmix_info_t` containing node-realm information. At a minimum,
10 either the `PMIX_NODEID` or `PMIX_HOSTNAME` attribute is required to be included in the
11 array, though both may be included.

12 Note that these assemblages can be used hierarchically:

- 13 • a `PMIX_JOB_INFO_ARRAY` might contain multiple `PMIX_APP_INFO_ARRAY` elements,
14 each describing values for a specific application within the job.
- 15 • a `PMIX_JOB_INFO_ARRAY` could contain a `PMIX_NODE_INFO_ARRAY` for each node
16 hosting processes from that job, each array describing job-level values for that node.
- 17 • a `PMIX_SESSION_INFO_ARRAY` might contain multiple `PMIX_JOB_INFO_ARRAY`
18 elements, each describing a job executing within the session. Each job array could, in turn,
19 contain both application and node arrays, thus providing a complete picture of the active
20 operations within the allocation.

▼ Advice to PMIx library implementers ▼

21 PMIx implementations must be capable of properly parsing and storing any hierarchical depth of
22 information arrays. The resulting stored values are must to be accessible via both `PMIx_Get` and
23 `PMIx_Query_info_nb` APIs, assuming appropriate directives are provided by the caller.

24 17.2.3.2 Assembling the registration information

25 The following description is not intended to represent the actual layout of information in a given
26 PMIx library. Instead, it describes how information provided in the `info` parameter of the
27 `PMIx_server_register_namespace` shall be organized for proper processing by a PMIx server
28 library. The ordering of the various information elements is arbitrary - they are presented in a
29 top-down hierarchical form solely for clarity in reading.

Advice to PMIx server hosts

1 Creating the *info* array of data requires knowing in advance the number of elements required for the
2 array. This can be difficult to compute and somewhat fragile in practice. One method for resolving
3 the problem is to create a linked list of objects, each containing a single `pmix_info_t` structure.
4 Allocation and manipulation of the list can then be accomplished using existing standard methods.
5 Upon completion, the final *info* array can be allocated based on the number of elements on the list,
6 and then the values in the list object `pmix_info_t` structures transferred to the corresponding
7 array element utilizing the `PMIx_Info_xfer` API.

8 A common building block used in several areas is the construction of a regular expression
9 identifying the nodes involved in that area - e.g., the nodes in a *session* or *job*. PMIx provides
10 several tools to facilitate this operation, beginning by constructing an argv-like array of node
11 names. This array is then passed to the `PMIx_generate_regex` function to create a regular
12 expression parseable by the PMIx server library, as shown below:

```
13 char **nodes = NULL;  
14 char *nodelist;  
15 char *regex;  
16 size_t n;  
17 pmix_status_t rc;  
18 pmix_info_t info;  
19  
20 /* loop over an array of nodes, adding each  
21  * name to the array */  
22 for (n=0; n < num_nodes; n++) {  
23     /* filter the nodes to ignore those not included  
24     * in the target range (session, job, etc.). In  
25     * this example, all nodes are accepted */  
26     PMIX_ARGV_APPEND(&nodes, node[n]->name);  
27 }  
28  
29 /* join into a comma-delimited string */  
30 nodelist = PMIX_ARGV_JOIN(nodes, ',');  
31  
32 /* release the array */  
33 PMIX_ARGV_FREE(nodes);  
34  
35 /* generate regex */  
36 rc = PMIx_generate_regex(nodelist, &regex);  
37  
38 /* release list */
```

```

1     free(nodelist);
2
3     /* pass the regex as the value to the PMIX_NODE_MAP key */
4     PMIx_Info_load(&info, PMIX_NODE_MAP, regex, PMIX_REGEX);
5     /* release the regex */
6     free(regex);

```

C

7 Changing the filter criteria allows the construction of node maps for any level of information. A
8 description of the returned regular expression is provided [here](#).

9 A similar method is used to construct the map of processes on each node from the namespace being
10 registered. This may be done for each information level of interest (e.g., to identify the process map
11 for the entire *job* or for each *application* in the job) by changing the search criteria. An example is
12 shown below for the case of creating the process map for a *job*:

C

```

13 char **ndppn;
14 char rank[30];
15 char **ppnarray = NULL;
16 char *ppn;
17 char *localranks;
18 char *regex;
19 size_t n, m;
20 pmix_status_t rc;
21 pmix_info_t info;
22
23 /* loop over an array of nodes */
24 for (n=0; n < num_nodes; n++) {
25     /* for each node, construct an array of ranks on that node */
26     ndppn = NULL;
27     for (m=0; m < node[n]->num_procs; m++) {
28         /* ignore processes that are not part of the target job */
29         if (!PMIX_CHECK_NAMESPACE(target_job, node[n]->proc[m].namespace)) {
30             continue;
31         }
32         snprintf(rank, 30, "%d", node[n]->proc[m].rank);
33         PMIX_ARGV_APPEND(&ndppn, rank);
34     }
35     /* convert the array into a comma-delimited string of ranks */
36     localranks = PMIX_ARGV_JOIN(ndppn, ',');
37     /* release the local array */
38     PMIX_ARGV_FREE(ndppn);
39     /* add this node's contribution to the overall array */

```

```

1     PMIX_ARGV_APPEND(&ppnarray, localranks);
2     /* release the local list */
3     free(localranks);
4 }
5
6 /* join into a semicolon-delimited string */
7 ppn = PMIX_ARGV_JOIN(ppnarray, ';');
8
9 /* release the array */
10 PMIX_ARGV_FREE(ppnarray);
11
12 /* generate ppn regex */
13 rc = PMIx_generate_ppn(ppn, &regex);
14
15 /* release list */
16 free(ppn);
17
18 /* pass the regex as the value to the PMIX_PROC_MAP key */
19 PMIx_Info_load(&info, PMIX_PROC_MAP, regex, PMIX_REGEX);
20 /* release the regex */
21 free(regex);

```

C

Note that the **PMIX_NODE_MAP** and **PMIX_PROC_MAP** attributes are linked in that the order of entries in the process map must match the ordering of nodes in the node map - i.e., there is no provision in the PMIx process map regular expression generator/parser pair supporting an out-of-order node or a node that has no corresponding process map entry (e.g., a node with no processes on it). Armed with these tools, the registration *info* array can be constructed as follows:

- Session-level information includes all session-specific values. In many cases, only two values (**PMIX_SESSION_ID** and **PMIX_UNIV_SIZE**) are included in the registration array. Since both of these values are session-specific, they can be specified independently - i.e., in their own **pmix_info_t** elements of the *info* array. Alternatively, they can be provided as a **pmix_data_array_t** array of **pmix_info_t** using the **PMIX_SESSION_INFO_ARRAY** attribute and identified by including the **PMIX_SESSION_ID** attribute in the array - this is required in cases where non-specific attributes (e.g., **PMIX_NUM_NODES** or **PMIX_NODE_MAP**) are passed to describe aspects of the session. Note that the node map can include nodes not used by the job being registered as no corresponding process map is specified.

The *info* array at this point might look like (where the labels identify the corresponding attribute - e.g., “Session ID” corresponds to the **PMIX_SESSION_ID** attribute):

- Job-level information includes all job-specific values such as **PMIX_JOB_SIZE**, **PMIX_JOB_NUM_APPS**, and **PMIX_JOBID**. Since each invocation of **PMIx_server_register_nspace** describes a single *job*, job-specific values can be




Figure 17.1.: Session-level information elements

1 specified independently - i.e., in their own `pmix_info_t` elements of the `info` array.

2 Alternatively, they can be provided as a `pmix_data_array_t` array of `pmix_info_t`

3 identified by the `PMIX_JOB_INFO_ARRAY` attribute - this is required in cases where

4 non-specific attributes (e.g., `PMIX_NODE_MAP`) are passed to describe aspects of the job. Note

5 that since the invocation only involves a single namespace, there is no need to include the

6 `PMIX_NAMESPACE` attribute in the array.

7 Upon conclusion of this step, the `info` array might look like:

8 Note that in this example, `PMIX_NUM_NODES` is not required as that information is contained in

9 the `PMIX_NODE_MAP` attribute. Similarly, `PMIX_JOB_SIZE` is not technically required as that

10 information is contained in the `PMIX_PROC_MAP` when combined with the corresponding node

11 map - however, there is no issue with including the job size as a separate entry.

12 The example also illustrates the hierarchical use of the `PMIX_NODE_INFO_ARRAY` attribute.

13 In this case, we have chosen to pass several job-related values for each node - since those values

14 are non-unique across the job, they must be passed in a node-info container. Note that the choice

15 of what information to pass into the PMIx server library versus what information to derive from

16 other values at time of request is left to the host environment. PMIx implementors in turn may, if

17 they choose, pre-parse registration data to create expanded views (thus enabling faster response

18 to requests at the expense of memory footprint) or to compress views into tighter representations

19 (thus trading minimized footprint for longer response times).

- 20 • Application-level information includes all application-specific values such as `PMIX_APP_SIZE`
 21 and `PMIX_APPLDR`. If the `job` contains only a single `application`, then the application-specific
 22 values can be specified independently - i.e., in their own `pmix_info_t` elements of the `info`
 23 array - or as a `pmix_data_array_t` array of `pmix_info_t` using the
 24 `PMIX_APP_INFO_ARRAY` attribute and identified by including the `PMIX_APPNUM` attribute in
 25 the array. Use of the array format is must in cases where non-specific attributes (e.g.,
 26 `PMIX_NODE_MAP`) are passed to describe aspects of the application.

27 However, in the case of a job consisting of multiple applications, all application-specific values

28 for each application must be provided using the `PMIX_APP_INFO_ARRAY` format, each

29 identified by its `PMIX_APPNUM` value.




Figure 17.2.: Job-level information elements

1 Upon conclusion of this step, the *info* array might look like that shown in 17.3, assuming there
 2 are two applications in the job being registered:

- 3 • Process-level information includes an entry for each process in the job being registered, each
 4 entry marked with the `PMIX_PROC_INFO_ARRAY` attribute. The *rank* of the process must be
 5 the first entry in the array - this provides efficiency when storing the data. Upon conclusion of
 6 this step, the *info* array might look like the diagram in 17.4:
- 7 • For purposes of this example, node-level information only includes values describing the local
 8 node - i.e., it does not include information about other nodes in the job or session. In many cases,
 9 the values included in this level are unique to it and can be specified independently - i.e., in their
 10 own `pmix_info_t` elements of the *info* array. Alternatively, they can be provided as a
 11 `pmix_data_array_t` array of `pmix_info_t` using the `PMIX_NODE_INFO_ARRAY`
 12 attribute - this is required in cases where non-specific attributes are passed to describe aspects of
 13 the node, or where values for multiple nodes are being provided.

14 The node-level information requires two elements that must be constructed in a manner similar to
 15 that used for the node map. The `PMIX_LOCAL_PEERS` value is computed based on the
 16 processes on the local node, filtered to select those from the job being registered, as shown below




Figure 17.3.: Application-level information elements

1 using the tools provided by PMIx:

```

2 char **ndppn = NULL;
3 char rank[30];
4 char *localranks;
5 size_t m;
6 pmix_info_t info;
7
8 for (m=0; m < mynode->num_procs; m++) {
9     /* ignore processes that are not part of the target job */
10    if (!PMIX_CHECK_NAMESPACE(targetjob,mynode->proc[m].nspace)) {
11        continue;
12    }
13    snprintf(rank, 30, "%d", mynode->proc[m].rank);
14    PMIX_ARGV_APPEND(&ndppn, rank);
15 }
16 /* convert the array into a comma-delimited string of ranks */
17 localranks = PMIX_ARGV_JOIN(ndppn, ',');
18 /* release the local array */
19 PMIX_ARGV_FREE(ndppn);

```




Figure 17.4.: Process-level information elements

```

1
2      /* pass the string as the value to the PMIX_LOCAL_PEERS key */
3      PMIx_Info_load(&info, PMIX_LOCAL_PEERS, localranks, PMIX_STRING);
4
5      /* release the list */
6      free(localranks);

```

C

7 The **PMIX_LOCAL_CPUSSETS** value is constructed in a similar manner. In the provided
8 example, it is assumed that an Hardware Locality (HWLOC) cpuset representation (a
9 comma-delimited string of processor IDs) of the processors assigned to each process has
10 previously been generated and stored on the process description. Thus, the value can be
11 constructed as shown below:

```

1     char **ndcpus = NULL;
2     char *localcpus;
3     size_t m;
4     pmix_info_t info;
5
6     for (m=0; m < mynode->num_procs; m++) {
7         /* ignore processes that are not part of the target job */
8         if (!PMIX_CHECK_NAMESPACE(targetjob, mynode->proc[m].nspace)) {
9             continue;
10        }
11        PMIX_ARGV_APPEND(&ndcpus, mynode->proc[m].cpuset);
12    }
13    /* convert the array into a colon-delimited string */
14    localcpus = PMIX_ARGV_JOIN(ndcpus, ':');
15    /* release the local array */
16    PMIX_ARGV_FREE(ndcpus);
17
18    /* pass the string as the value to the PMIX_LOCAL_CPUSSETS key */
19    PMIX_Info_load(&info, PMIX_LOCAL_CPUSSETS, localcpus, PMIX_STRING);
20
21    /* release the list */
22    free(localcpus);

```

23 Note that for efficiency, these two values can be computed at the same time.

24 The final *info* array might therefore look like the diagram in 17.5:

25 17.2.4 PMIx_server_deregister_namespace

26 Summary

27 Deregister a namespace.

28 Format

PMIx v1.0

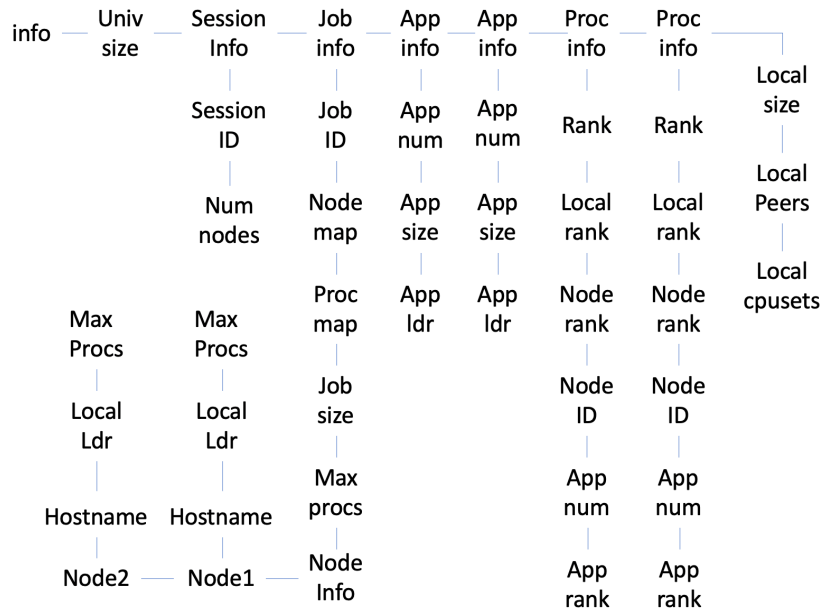


Figure 17.5.: Final information array

```

1  void PMIx_server_deregister_nspace(const pmix_namespace_t nspace,
2  pmix_op_cbfunc_t cbfunc, void *cbdata);

```

- 3 **IN nspace**
Namespace (string)
- 4 **IN cbfunc**
Callback function `pmix_op_cbfunc_t`. A **NULL** function reference indicates that the
- 5 function is to be executed as a blocking operation. (function reference)
- 6 **IN cbdata**
Data to be passed to the callback function (memory reference)

10 **Description**

11 Deregister the specified *nspace* and purge all objects relating to it, including any client information
12 from that namespace. This is intended to support persistent PMIx servers by providing an
13 opportunity for the host RM to tell the PMIx server library to release all memory for a completed
14 job. Note that the library must not invoke the callback function prior to returning from the API, and
15 that a **NULL** *cbfunc* reference indicates that the function is to be executed as a blocking operation.

1 17.2.5 PMIx_server_register_resources

2 Summary

3 Register non-namespace related information with the local PMIx server library.

4 *PMIx v4.0* Format

C

5 `pmix_status_t`

```
6 PMIx_server_register_resources(pmix_info_t info[], size_t ninfo,  
7                              pmix_op_cbfunc_t cbfunc,  
8                              void *cbdata);
```

C

9 **IN** `info`

10 Array of info structures (array of handles)

11 **IN** `ninfo`

12 Number of elements in the *info* array (integer)

13 **IN** `cbfunc`

14 Callback function `pmix_op_cbfunc_t`. A **NULL** function reference indicates that the
15 function is to be executed as a blocking operation (function reference)

16 **IN** `cbdata`

17 Data to be passed to the callback function (memory reference)

18 Description

19 Pass information about resources not associated with a given namespace to the PMIx server library
20 for distribution to local client processes. This includes information on fabric devices, GPUs, and
21 other resources. All information provided through this API shall be made available to each job as
22 part of its job-level information. Duplicate information provided with the
23 `PMIx_server_register_namespace` API shall override any information provided by this
24 function for that namespace, but only for that specific namespace.

25 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Advice to PMIx server hosts

26 Note that information passed in this manner could also have been included in a call to
27 `PMIx_server_register_namespace` - e.g., as part of a `PMIX_NODE_INFO_ARRAY` array.
28 This API is provided as a logical alternative for code clarity, especially where multiple jobs may be
29 supported by a single PMIx server library instance, to avoid multiple registration of static resource
30 information.

31 A **NULL** *cbfunc* reference indicates that the function is to be executed as a blocking operation.

1 17.2.6 PMIx_server_deregister_resources

2 Summary

3 Remove specified non-namespace related information from the local PMIx server library.

4 Format

PMIx v4.0

```
5 pmix_status_t
6 PMIx_server_deregister_resources(pmix_info_t info[], size_t ninfo,
7                                 pmix_op_cbfunc_t cbfunc,
8                                 void *cbdata);
```

9 IN info

10 Array of info structures (array of handles)

11 IN ninfo

12 Number of elements in the *info* array (integer)

13 IN cbfunc

14 Callback function [pmix_op_cbfunc_t](#). A **NULL** function reference indicates that the
15 function is to be executed as a blocking operation (function reference)

16 IN cbdata

17 Data to be passed to the callback function (memory reference)

18 Description

19 Remove information about resources not associated with a given namespace from the PMIx server
20 library. Only the *key* fields of the provided *info* array shall be used for the operation - the associated
21 values shall be ignored except where they serve as qualifiers to the request. For example, to remove
22 a specific fabric device from a given node, the *info* array might include a
23 [PMIX_NODE_INFO_ARRAY](#) containing the [PMIX_NODEID](#) or [PMIX_HOSTNAME](#) identifying
24 the node hosting the device, and the [PMIX_FABRIC_DEVICE_NAME](#) specifying the device to be
25 removed. Alternatively, the device could be removed using only the [PMIX_DEVICE_ID](#) as this is
26 unique across the overall system.

27 Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Advice to PMIx server hosts

28 As information not related to namespaces is considered *static*, there is no requirement that the host
29 environment deregister resources prior to finalizing the PMIx server library. The server library
30 shall properly cleanup as part of its normal finalize operations. Deregistration of resources is only
31 required, therefore, when the host environment determines that client processes should no longer
32 have access to that information.

33 A **NULL** *cbfunc* reference indicates that the function is to be executed as a blocking operation.

1 17.2.7 PMIx_server_register_client

2 Summary

3 Register a client process with the PMIx server library.

4 *PMIx v1.0* Format

C

```
5 pmix_status_t  
6 PMIx_server_register_client(const pmix_proc_t *proc,  
7                             uid_t uid, gid_t gid,  
8                             void *server_object,  
9                             pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

10 **IN** `proc`
11 `pmix_proc_t` structure (handle)
12 **IN** `uid`
13 user id (integer)
14 **IN** `gid`
15 group id (integer)
16 **IN** `server_object`
17 (memory reference)
18 **IN** `cbfunc`
19 Callback function `pmix_op_cbfunc_t`. A **NULL** function reference indicates that the
20 function is to be executed as a blocking operation (function reference)
21 **IN** `cbdata`
22 Data to be passed to the callback function (memory reference)

23 A successful return indicates that the request is being processed and the result will be returned in
24 the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning
25 from the API. The callback function, `cbfunc`, is only called when **PMIX_SUCCESS** is returned.

26 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 27 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
28 returned *success* - the `cbfunc` will not be called

29 If none of the above return codes are appropriate, then an implementation must return either a
30 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Description

Register a client process with the PMIx server library.

The host server can also, if it desires, provide an object it wishes to be returned when a server function is called that relates to a specific process. For example, the host server may have an object that tracks the specific client. Passing the object to the library allows the library to provide that object to the host server during subsequent calls related to that client, such as a `pmix_server_client_connected2_fn_t` function. This allows the host server to access the object without performing a lookup based on the client's namespace and rank.

Advice to PMIx server hosts

Host environments are required to execute this operation prior to starting the client process. The expected user ID and group ID of the child process allows the server library to properly authenticate clients as they connect by requiring the two values to match. Accordingly, the detected user and group ID's of the connecting process are not included in the `pmix_server_client_connected2_fn_t` server module function.

Advice to PMIx library implementers

For security purposes, the PMIx server library should check the user and group ID's of a connecting process against those provided for the declared client process identifier via the `PMIx_server_register_client` prior to completing the connection.

17.2.8 `PMIx_server_deregister_client`

Summary

Deregister a client and purge all data relating to it.

Format

PMIx v1.0

```
void  
PMIx_server_deregister_client(const pmix_proc_t *proc,  
                             pmix_op_cbfunc_t cbfunc, void *cbdata);
```

IN `proc`
`pmix_proc_t` structure (handle)

IN `cbfunc`
Callback function `pmix_op_cbfunc_t`. A **NULL** function reference indicates that the function is to be executed as a blocking operation (function reference)

IN `cbdata`
Data to be passed to the callback function (memory reference)

Description

The `PMIx_server_deregister_nspace` API will delete all client information for that namespace. The PMIx server library will automatically perform that operation upon disconnect of all local clients. This API is therefore intended primarily for use in exception cases, but can be called in non-exception cases if desired. Note that the library must not invoke the callback function prior to returning from the API.

17.2.9 PMIx_server_setup_fork

Summary

Setup the environment of a child process to be forked by the host.

Format

PMIx v1.0

```
pmix_status_t
PMIx_server_setup_fork(const pmix_proc_t *proc,
                      char ***env);
```

IN proc

`pmix_proc_t` structure (handle)

IN env

Environment array (array of strings)

Returns `PMIX_SUCCESS` or a negative value indicating the error.

Description

Setup the environment of a child process to be forked by the host so it can correctly interact with the PMIx server.

The PMIx client needs some setup information so it can properly connect back to the server. This function will set appropriate environmental variables for this purpose, and will also provide any environmental variables that were specified in the launch command (e.g., via `PMIx_Spawn`) plus other values (e.g., variables required to properly initialize the client's fabric library).

Advice to PMIx server hosts

Host environments are required to execute this operation prior to starting the client process.

17.2.10 PMIx_server_dmodex_request

Summary

Define a function by which the host server can request modex data from the local PMIx server.

Format

C

```
pmix_status_t
PMIx_server_dmodex_request(const pmix_proc_t *proc,
                           pmix_dmodex_response_fn_t cbfunc,
                           void *cbdata);
```

C

IN `proc`

`pmix_proc_t` structure (handle)

IN `cbfunc`

Callback function `pmix_dmodex_response_fn_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning from the API. The callback function, `cbfunc`, is only called when `PMIX_SUCCESS` is returned.

Description

Define a function by which the host server can request modex data from the local PMIx server. Traditional wireup procedures revolve around the per-process posting of data (e.g., location and endpoint information) via the `PMIx_Put` and `PMIx_Commit` functions followed by a `PMIx_Fence` barrier that globally exchanges the posted information. However, the barrier operation represents a significant time impact at large scale.

PMIx supports an alternative wireup method known as *Direct Modex* that replaces the barrier-based exchange of all process-posted information with on-demand fetch of a peer's data. In place of the barrier operation, data posted by each process is cached on the local PMIx server. When a process requests the information posted by a particular peer, it first checks the local cache to see if the data is already available. If not, then the request is passed to the local PMIx server, which subsequently requests that its RM host request the data from the RM daemon on the node where the specified peer process is located. Upon receiving the request, the RM daemon passes the request into its PMIx server library using the `PMIx_server_dmodex_request` function, receiving the response in the provided `cbfunc` once the indicated process has posted its information. The RM daemon then returns the data to the requesting daemon, who subsequently passes the data to its PMIx server library for transfer to the requesting client.

Advice to users

While direct modex allows for faster launch times by eliminating the barrier operation, per-peer retrieval of posted information is less efficient. Optimizations can be implemented - e.g., by returning posted information from all processes on a node upon first request - but in general direct modex remains best suited for sparsely connected applications.

1 17.2.10.1 Server Direct Modex Response Callback Function

2 The `PMIx_server_dmodex_request` callback function.

3 Summary

4 Provide a function by which the local PMIx server library can return connection and other data
5 posted by local application processes to the host resource manager.

6 Format

PMIx v1.0

C

```
7 typedef void (*pmix_dmodex_response_fn_t) (  
8             pmix_status_t status,  
9             char *data, size_t sz,  
10            void *cbdata);
```

C

11 IN status

12 Returned status of the request (`pmix_status_t`)

13 IN data

14 Pointer to a data "blob" containing the requested information (handle)

15 IN sz

16 Number of bytes in the *data* blob (integer)

17 IN cbdata

18 Data passed into the initial call to `PMIx_server_dmodex_request` (memory reference)

19 Description

20 Define a function to be called by the PMIx server library for return of information posted by a local
21 application process (via `PMIx_Put` with subsequent `PMIx_Commit`) in response to a request
22 from the host RM. The returned *data* blob is owned by the PMIx server library and will be free'd
23 upon return from the function.

24 17.2.11 PMIx_server_setup_application

25 Summary

26 Provide a function by which a launcher can request application-specific setup data prior to launch of
27 a *job*.

Format

C

```
pmix_status_t
PMIX_server_setup_application(const pmix_namespace_t nspace,
                             pmix_info_t info[], size_t ninfo,
                             pmix_setup_application_cbfunc_t cbfunc,
                             void *cbdata);
```

C

IN nspace
namespace (string)

IN info
Array of info structures (array of handles)

IN ninfo
Number of elements in the *info* array (integer)

IN cbfunc
Callback function `pmix_setup_application_cbfunc_t` (function reference)

IN cbdata
Data to be passed to the *cbfunc* callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Required Attributes

PMIx libraries that support this operation are required to support the following:

PMIX_SETUP_APP_ENVARS "pmix.setup.env" (bool)
Harvest and include relevant environmental variables.

PMIX_SETUP_APP_NONENVARS "pmix.setup.nenv" (bool)
Include all relevant data other than environmental variables.

PMIX_SETUP_APP_ALL "pmix.setup.all" (bool)
Include all relevant data.

PMIX_ALLOC_FABRIC "pmix.alloc.net" (array)
Array of `pmix_info_t` describing requested fabric resources. This must include at least: **PMIX_ALLOC_FABRIC_ID**, **PMIX_ALLOC_FABRIC_TYPE**, and **PMIX_ALLOC_FABRIC_ENDPTS**, plus whatever other descriptors are desired.

PMIX_ALLOC_FABRIC_ID "pmix.alloc.netid" (char*)

1 The key to be used when accessing this requested fabric allocation. The fabric allocation
2 will be returned/stored as a `pmix_data_array_t` of `pmix_info_t` whose first
3 element is composed of this key and the allocated resource description. The type of the
4 included value depends upon the fabric support. For example, a TCP allocation might
5 consist of a comma-delimited string of socket ranges such as "32000–32100,
6 33005, 38123–38146". Additional array entries will consist of any provided resource
7 request directives, along with their assigned values. Examples include:
8 `PMIX_ALLOC_FABRIC_TYPE` - the type of resources provided;
9 `PMIX_ALLOC_FABRIC_PLANE` - if applicable, what plane the resources were assigned
10 from; `PMIX_ALLOC_FABRIC_QOS` - the assigned QoS; `PMIX_ALLOC_BANDWIDTH` -
11 the allocated bandwidth; `PMIX_ALLOC_FABRIC_SEC_KEY` - a security key for the
12 requested fabric allocation. NOTE: the array contents may differ from those requested,
13 especially if `PMIX_INFO_REQD` was not set in the request.

14 `PMIX_ALLOC_FABRIC_SEC_KEY` "pmix.alloc.nsec" (`pmix_byte_object_t`)
15 Request that the allocation include a fabric security key for the spawned job.

16 `PMIX_ALLOC_FABRIC_TYPE` "pmix.alloc.nettype" (`char*`)
17 Type of desired transport (e.g., "tcp", "udp") being requested in an allocation request.

18 `PMIX_ALLOC_FABRIC_PLANE` "pmix.alloc.netplane" (`char*`)
19 ID string for the *fabric plane* to be used for the requested allocation.

20 `PMIX_ALLOC_FABRIC_ENDPTS` "pmix.alloc.endpts" (`size_t`)
21 Number of endpoints to allocate per *process* in the job.

22 `PMIX_ALLOC_FABRIC_ENDPTS_NODE` "pmix.alloc.endpts.nd" (`size_t`)
23 Number of endpoints to allocate per *node* for the job.

24 `PMIX_PROC_MAP` "pmix.pmap" (`char*`)
25 Regular expression describing processes on each node in the specified realm - see 17.2.3.2
26 for an explanation of its generation. Defaults to the *job* realm.

27 `PMIX_NODE_MAP` "pmix.nmap" (`char*`)
28 Regular expression of nodes currently hosting processes in the specified realm - see 17.2.3.2
29 for an explanation of its generation. Defaults to the *job* realm.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

30 PMIx libraries that support this operation may support the following:

31 `PMIX_ALLOC_BANDWIDTH` "pmix.alloc.bw" (`float`)
32 Fabric bandwidth (in Megabits[base2]/sec) for the job being requested in an allocation
33 request.

34 `PMIX_ALLOC_FABRIC_QOS` "pmix.alloc.netqos" (`char*`)
35 Fabric quality of service level for the job being requested in an allocation request.

1 **PMIX_SESSION_INFO** "pmix.ssn.info" (bool)

2 Return information regarding the session realm of the target process. In this context,
3 indicates that the information provided in the **PMIX_NODE_MAP** is for the entire session and
4 not just the indicated namespace. Thus, subsequent calls to this API may omit node-level
5 information - e.g., the library may not need to include information on the devices on each
6 node in a subsequent call.

7 The following optional attributes may be provided by the host environment to identify the
8 programming model (as specified by the user) being executed within the application. The PMIx
9 server library may utilize this information to harvest/forward model-specific environmental
10 variables, record the programming model associated with the application, etc.

- 11 • **PMIX_PROGRAMMING_MODEL** "pmix.pgm.model" (char*)
12 Programming model being initialized (e.g., "MPI" or "OpenMP").
- 13 • **PMIX_MODEL_LIBRARY_NAME** "pmix.mdl.name" (char*)
14 Programming model implementation ID (e.g., "OpenMPI" or "MPICH").
- 15 • **PMIX_MODEL_LIBRARY_VERSION** "pmix.mld.vrs" (char*)
16 Programming model version string (e.g., "2.1.1").



17 Description

18 Provide a function by which the RM can request application-specific setup data (e.g., environmental
19 variables, fabric configuration and security credentials) from supporting PMIx server library
20 subsystems prior to initiating launch of a job.

21 This is defined as a non-blocking operation in case contributing subsystems need to perform some
22 potentially time consuming action (e.g., query a remote service) before responding. The returned
23 data must be distributed by the host environment and subsequently delivered to the local PMIx
24 server on each node where application processes will execute, prior to initiating execution of those
25 processes.

Advice to PMIx server hosts

26 Host environments are required to execute this operation prior to launching a job. In addition to
27 supported directives, the *info* array must include a description of the *job* using the
28 **PMIX_NODE_MAP** and **PMIX_PROC_MAP** attributes.

29 Note that the function can be called on a per-application basis if the **PMIX_PROC_MAP** and
30 **PMIX_NODE_MAP** are provided only for the corresponding application (as opposed to the entire
31 job) each time.

Advice to PMIx library implementers

32 Support for harvesting of environmental variables and providing of local configuration information
33 by the PMIx implementation is optional.

1 17.2.11.1 Server Setup Application Callback Function

2 The `PMIx_server_setup_application` callback function.

3 Summary

4 Provide a function by which the resource manager can receive application-specific environmental
5 variables and other setup data prior to launch of an application.

6 Format

PMIx v2.0

C

```
typedef void (*pmix_setup_application_cfunc_t) (  
    pmix_status_t status,  
    pmix_info_t info[], size_t ninfo,  
    void *provided_cbdata,  
    pmix_op_cfunc_t cfunc, void *cbdata);
```

C

12 IN status

13 returned status of the request (`pmix_status_t`)

14 IN info

15 Array of info structures (array of handles)

16 IN ninfo

17 Number of elements in the *info* array (integer)

18 IN provided_cbdata

19 Data originally passed to call to `PMIx_server_setup_application` (memory
20 reference)

21 IN cfunc

22 `pmix_op_cfunc_t` function to be called when processing completed (function reference)

23 IN cbdata

24 Data to be passed to the *cbfunc* callback function (memory reference)

25 Description

26 Define a function to be called by the PMIx server library for return of application-specific setup
27 data in response to a request from the host RM. The returned *info* array is owned by the PMIx
28 server library and will be free'd when the provided *cbfunc* is called.

29 17.2.11.2 Server Setup Application Attributes

PMIx v3.0

30 Attributes specifically defined for controlling contents of application setup data.

31 `PMIX_SETUP_APP_ENVARS "pmix.setup.env" (bool)`

32 Harvest and include relevant environmental variables.

33 `PMIX_SETUP_APP_NONENVARS ""pmix.setup.nenv" (bool)`

34 Include all relevant data other than environmental variables.

35 `PMIX_SETUP_APP_ALL "pmix.setup.all" (bool)`

36 Include all relevant data.

1 17.2.12 `PMIx_Register_attributes`

2 Summary

3 Register host environment attribute support for a function.

4 *PMIx v4.0* Format

C

5 `pmix_status_t`

```
6 PMIx_Register_attributes(char *function,  
7                         pmix_regattr_t attrs[],  
8                         size_t nattrs);
```

C

9 **IN** `function`

String name of function (string)

11 **IN** `attrs`

Array of `pmix_regattr_t` describing the supported attributes (handle)

13 **IN** `nattrs`

Number of elements in `attrs` (`size_t`)

15 Returns `PMIX_SUCCESS` or a negative value indicating the error.

16 Description

17 The `PMIx_Register_attributes` function is used by the host environment to register with
18 its PMIx server library the attributes it supports for each `pmix_server_module_t` function.

19 The *function* is the string name of the server module function (e.g., "register_events",
20 "validate_credential", or "allocate") whose attributes are being registered. See the
21 `pmix_regattr_t` entry for a description of the *attrs* array elements.

22 Note that the host environment can also query the library (using the `PMIx_Query_info_nb`
23 API) for its attribute support both at the server, client, and tool levels once the host has executed
24 `PMIx_server_init` since the server will internally register those values.

Advice to PMIx server hosts

25 Host environments are strongly encouraged to register all supported attributes immediately after
26 initializing the library to ensure that user requests are correctly serviced.

Advice to PMIx library implementers

PMIx implementations are *required* to register all internally supported attributes for each API during initialization of the library (i.e., when the process calls their respective PMIx init function). Specifically, the implementation *must not* register supported attributes upon first call to a given API as this would prevent users from discovering supported attributes prior to first use of an API.

It is the implementation's responsibility to associate registered attributes for a given `pmix_server_module_t` function with their corresponding user-facing API. Supported attributes *must* be reported to users in terms of their support for user-facing APIs, broken down by the level (see Section 7.1.5) at which the attribute is supported.

Note that attributes can/will be registered on an API for each level. It is *required* that the implementation support user queries for supported attributes on a per-level basis. Duplicate registrations at the *same* level for a function *shall* return an error - however, duplicate registrations at *different* levels *shall* be independently tracked.

17.2.12.1 Attribute registration constants

Constants supporting attribute registration.

`PMIX_ERR_REPEAT_ATTR_REGISTRATION` -171 The attributes for an identical function have already been registered at the specified level (host, server, or client).

17.2.12.2 Attribute registration structure

The `pmix_regattr_t` structure is used to register attribute support for a PMIx function.

PMIx v4.0

```
typedef struct pmix_regattr {
    char *name;
    pmix_key_t *string;
    pmix_data_type_t type;
    pmix_info_t *info;
    size_t ninfo;
    char **description;
} pmix_regattr_t;
```

Note that in this structure:

- the *name* is the actual name of the attribute - e.g., "PMIX_MAX_PROCS"
- the *string* is the literal string value of the attribute - e.g., "pmix.max.size" for the `PMIX_MAX_PROCS` attribute
- *type* must be a PMIx data type identifying the type of data associated with this attribute.

- the *info* array contains machine-usable information regarding the range of accepted values. This may include entries for `PMIX_MIN_VALUE`, `PMIX_MAX_VALUE`, `PMIX_ENUM_VALUE`, or a combination of them. For example, an attribute that supports all positive integers might delineate it by including a `pmix_info_t` with a key of `PMIX_MIN_VALUE`, type of `PMIX_INT`, and value of zero. The lack of an entry for `PMIX_MAX_VALUE` indicates that there is no ceiling to the range of accepted values.
- *ninfo* indicates the number of elements in the *info* array
- The *description* field consists of a `NULL`-terminated array of strings describing the attribute, optionally including a human-readable description of the range of accepted values - e.g., "ALL POSITIVE INTEGERS", or a comma-delimited list of enum value names. No correlation between the number of entries in the *description* and the number of elements in the *info* array is implied or required.

The attribute *name* and *string* fields must be `NULL`-terminated strings composed of standard alphanumeric values supported by common utilities such as *strcmp*.

Although not strictly required, both PMIx library implementers and host environments are strongly encouraged to provide both human-readable and machine-parsable descriptions of supported attributes when registering them.

17.2.12.3 Attribute registration structure descriptive attributes

The following attributes relate to the nature of the values being reported in the `pmix_regattr_t` structures.

PMIX_MAX_VALUE "pmix.descr.maxval" (varies)

Used in `pmix_regattr_t` to describe the maximum valid value for the associated attribute.

PMIX_MIN_VALUE "pmix.descr.minval" (varies)

Used in `pmix_regattr_t` to describe the minimum valid value for the associated attribute.

PMIX_ENUM_VALUE "pmix.descr.enum" (char*)

Used in `pmix_regattr_t` to describe accepted values for the associated attribute.

Numerical values shall be presented in a form convertible to the attribute's declared data type. Named values (i.e., values defined by constant names via a typical C-language enum declaration) must be provided as their numerical equivalent.

17.2.12.4 Attribute registration structure support macros

The following macros are provided to support the `pmix_regattr_t` structure.

Static initializer for the regattr structure

Provide a static initializer for the `pmix_regattr_t` fields.

Provisional

PMIx v4.2

			C	
1		PMIX_REGATTR_STATIC_INIT		
			C	
2		Initialize the regattr structure		
3		Initialize the <code>pmix_regattr_t</code> fields		
	<i>PMIx v4.0</i>		C	
4		PMIX_REGATTR_CONSTRUCT (m)		
			C	
5		IN m		
6		Pointer to the structure to be initialized (pointer to <code>pmix_regattr_t</code>)		
7		Destruct the regattr structure		
8		Destruct the <code>pmix_regattr_t</code> fields, releasing all strings.		
	<i>PMIx v4.0</i>		C	
9		PMIX_REGATTR_DESTRUCT (m)		
			C	
10		IN m		
11		Pointer to the structure to be destructed (pointer to <code>pmix_regattr_t</code>)		
12		Create a regattr array		
13		Allocate and initialize an array of <code>pmix_regattr_t</code> structures.		
	<i>PMIx v4.0</i>		C	
14		PMIX_REGATTR_CREATE (m, n)		
			C	
15		INOUT m		
16		Address where the pointer to the array of <code>pmix_regattr_t</code> structures shall be stored		
17		(handle)		
18		IN n		
19		Number of structures to be allocated (size_t)		
20		Free a regattr array		
21		Release an array of <code>pmix_regattr_t</code> structures.		
	<i>PMIx v4.0</i>		C	
22		PMIX_REGATTR_FREE (m, n)		
			C	
23		INOUT m		
24		Pointer to the array of <code>pmix_regattr_t</code> structures (handle)		
25		IN n		
26		Number of structures in the array (size_t)		

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

Load a regattr structure

Load values into a `pmix_regattr_t` structure. The macro can be called multiple times to add as many strings as desired to the same structure by passing the same address and a **NULL** key to the macro. Note that the *t* type value must be given each time.

PMIx v4.0

```
▼ _____ C _____ ▼  
PMIX_REGATTR_LOAD(a, n, k, t, ni, v)  
▲ _____ C _____ ▲
```

- IN a**
Pointer to the structure to be loaded (pointer to `pmix_proc_t`)
- IN n**
String name of the attribute (string)
- IN k**
Key value to be loaded (`pmix_key_t`)
- IN t**
Type of data associated with the provided key (`pmix_data_type_t`)
- IN ni**
Number of `pmix_info_t` elements to be allocated in *info* (`size_t`)
- IN v**
One-line description to be loaded (more can be added separately) (string)

Transfer a regattr to another regattr

Non-destructively transfer the contents of a `pmix_regattr_t` structure to another one.

PMIx v4.0

```
▼ _____ C _____ ▼  
PMIX_REGATTR_XFER(m, n)  
▲ _____ C _____ ▲
```

- INOUT m**
Pointer to the destination `pmix_regattr_t` structure (handle)
- IN n**
Pointer to the source `pmix_regattr_t` structure (handle)

17.2.13 `PMIx_server_setup_local_support`

Summary

Provide a function by which the local `PMIx` server can perform any application-specific operations prior to spawning local clients of a given application.

Format

C

```
pmix_status_t
PMIx_server_setup_local_support(const pmix_namespace_t nspace,
                                pmix_info_t info[], size_t ninfo,
                                pmix_op_cbfunc_t cbfunc,
                                void *cbdata);
```

C

IN **nspace**

Namespace (string)

IN **info**

Array of info structures (array of handles)

IN **ninfo**

Number of elements in the *info* array (**size_t**)

IN **cbfunc**

Callback function **pmix_op_cbfunc_t**. A **NULL** function reference indicates that the function is to be executed as a blocking operation (function reference)

IN **cbdata**

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called

If none of the above return codes are appropriate, then an implementation must return either a general **PMIx** error code or an implementation defined error code as described in Section 3.1.1.

Description

Provide a function by which the local **PMIx** server can perform any application-specific operations prior to spawning local clients of a given application. For example, a fabric library might need to setup the local driver for “instant on” addressing. The data provided in the *info* array is the data returned to the host RM by the callback function executed as a result of a call to

PMIx_server_setup_application.

Advice to **PMIx** server hosts

Host environments are required to execute this operation prior to starting any local application processes from the specified namespace if information was obtained from a call to

PMIx_server_setup_application.

1 Host environments must register the *nspc* using `PMIx_server_register_nspace` prior to
2 calling this API to ensure that all namespace-related information required to support this function is
3 available to the library. This eliminates the need to include any of the registration information in the
4 *info* array passed to this API.

5 17.2.14 PMIx_server_IOF_deliver

6 Summary

7 Provide a function by which the host environment can pass forwarded Input/Output (IO) to the
8 PMIx server library for distribution to its clients.

9 Format

PMIx v3.0

```
10 pmix_status_t  
11 PMIx_server_IOF_deliver(const pmix_proc_t *source,  
12                        pmix_iof_channel_t channel,  
13                        const pmix_byte_object_t *bo,  
14                        const pmix_info_t info[], size_t ninfo,  
15                        pmix_op_cbfunc_t cbfunc, void *cbdata);
```

16 IN source

17 Pointer to `pmix_proc_t` identifying source of the IO (handle)

18 IN channel

19 IO channel of the data (`pmix_iof_channel_t`)

20 IN bo

21 Pointer to `pmix_byte_object_t` containing the payload to be delivered (handle)

22 IN info

23 Array of `pmix_info_t` metadata describing the data (array of handles)

24 IN ninfo

25 Number of elements in the *info* array (`size_t`)

26 IN cbfunc

27 Callback function `pmix_op_cbfunc_t`. A `NULL` function reference indicates that the
28 function is to be executed as a blocking operation (function reference)

29 IN cbdata

30 Data to be passed to the callback function (memory reference)

31 A successful return indicates that the request is being processed and the result will be returned in
32 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
33 from the API. The callback function, *cbfunc*, is only called when `PMIX_SUCCESS` is returned.

34 Returns `PMIX_SUCCESS` or one of the following error codes when the condition described occurs:

- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called

If none of the above return codes are appropriate, then an implementation must return either a general PMIx error code or an implementation defined error code as described in Section 3.1.1.

Description

Provide a function by which the host environment can pass forwarded IO to the PMIx server library for distribution to its clients. The PMIx server library is responsible for determining which of its clients have actually registered for the provided data and delivering it. The *cbfunc* callback function will be called once the PMIx server library no longer requires access to the provided data.

17.2.15 PMIx_server_collect_inventory

Summary

Collect inventory of resources on a node.

Format

PMIx v3.0

```
pmix_status_t
PMIx_server_collect_inventory(const pmix_info_t directives[],
                             size_t ndirs,
                             pmix_info_cbfunc_t cbfunc,
                             void *cbdata);
```

IN directives

Array of [pmix_info_t](#) directing the request (array of handles)

IN ndirs

Number of elements in the *directives* array ([size_t](#))

IN cbfunc

Callback function to return collected data ([pmix_info_cbfunc_t](#) function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

A successful return indicates that the request is being processed and the result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

Description

Provide a function by which the host environment can request its PMIx server library collect an inventory of local resources. Supported resources depends upon the PMIx implementation, but may include the local node topology and fabric interfaces.

Advice to PMIx server hosts

This is a non-blocking API as it may involve somewhat lengthy operations to obtain the requested information. Inventory collection is expected to be a rare event – at system startup and upon command from a system administrator. Inventory updates are expected to initiate a smaller operation involving only the changed information. For example, replacement of a node would generate an event to notify the scheduler with an inventory update without invoking a global inventory operation.

17.2.16 PMIx_server_deliver_inventory

Summary

Pass collected inventory to the PMIx server library for storage.

Format

PMIx v3.0

C

```
pmix_status_t
PMIx_server_deliver_inventory(const pmix_info_t info[],
                              size_t ninfo,
                              const pmix_info_t directives[],
                              size_t ndirs,
                              pmix_op_cbfunc_t cbfunc,
                              void *cbdata);
```

C

- IN info**
Array of `pmix_info_t` containing the inventory (array of handles)
- IN ninfo**
Number of elements in the *info* array (`size_t`)
- IN directives**
Array of `pmix_info_t` directing the request (array of handles)
- IN ndirs**
Number of elements in the *directives* array (`size_t`)
- IN cbfunc**
Callback function `pmix_op_cbfunc_t`. A **NULL** function reference indicates that the function is to be executed as a blocking operation (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

1 Returns one of the following:

2 A successful return indicates that the request is being processed and the result will be returned in
 3 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
 4 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

5 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

- 6 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
 7 returned *success* - the *cbfunc* will not be called

8 If none of the above return codes are appropriate, then an implementation must return either a
 9 general **PMIX** error code or an implementation defined error code as described in Section 3.1.1.

10 Description

11 Provide a function by which the host environment can pass inventory information obtained from a
 12 node (as a result of a call to **PMIx_server_collect_inventory**) to the **PMIx** server library
 13 for storage. Inventory data is subsequently used by the **PMIx** server library for allocations in
 14 response to **PMIx_server_setup_application**, and may be available to the library's host
 15 via the **PMIx_Get** API (depending upon **PMIx** implementation). The *cbfunc* callback function
 16 will be called once the **PMIx** server library no longer requires access to the provided data.

17 17.2.17 **PMIx_server_generate_locality_string**

18 Summary

19 Generate a **PMIx** locality string from a given *cpuset*.

20 Format

PMIx v4.0

C

21 **pmix_status_t**

22 **PMIx_server_generate_locality_string**(const **pmix_cpuset_t** **cpuset*,
 23 **char** ***locality*);

C

24 **IN** *cpuset*

25 Pointer to a **pmix_cpuset_t** containing the bitmap of assigned PUs (handle)

26 **OUT** *locality*

27 String representation of the **PMIx** locality corresponding to the input bitmap (**char***)

28 A successful return indicates that the returned string contains the generated locality string.

29 Returns **PMIX_SUCCESS** or a negative value indicating the error.

Description

Provide a function by which the host environment can generate a PMIx locality string for inclusion in the call to `PMIx_server_register_namespace`. This function shall only be called for local client processes, with the returned locality included in the job-level information (via the `PMIX_LOCALITY_STRING` attribute) provided to local clients. Local clients can use these strings as input to determine the relative locality of their local peers via the `PMIx_Get_relative_locality` API.

The function is required to return a string prefixed by the *source* field of the provided *cpuset* followed by a colon. The remainder of the string shall represent the corresponding locality as expressed by the underlying implementation.

17.2.18 `PMIx_server_generate_cpuset_string`

Summary

Generate a PMIx string representation of the provided *cpuset*.

Format

PMIx v4.0

```
pmix_status_t
PMIx_server_generate_cpuset_string(const pmix_cpuset_t *cpuset,
                                   char **cpuset_string);
```

IN *cpuset*

Pointer to a `pmix_cpuset_t` containing the bitmap of assigned PUs (handle)

OUT *cpuset_string*

String representation of the input bitmap (`char*`)

A successful return indicates that the returned string contains the generated *cpuset* representation string.

Returns `PMIX_SUCCESS` or a negative value indicating the error.

Description

Provide a function by which the host environment can generate a string representation of the *cpuset* bitmap for inclusion in the call to `PMIx_server_register_namespace`. This function shall only be called for local client processes, with the returned string included in the job-level information (via the `PMIX_CPUSSET` attribute) provided to local clients. Local clients can use these strings as input to obtain their PU bindings via the `PMIx_Parse_cpuset_string` API.

The function is required to return a string prefixed by the *source* field of the provided *cpuset* followed by a colon. The remainder of the string shall represent the PUs to which the process is bound as expressed by the underlying implementation.

1 17.2.18.1 Cpuset Structure

2 The `pmix_cpuset_t` structure contains a character string identifying the source of the bitmap
3 (e.g., "hwloc") and a pointer to the corresponding implementation-specific structure (e.g.,
4 `hwloc_cpuset_t`).

PMIx v4.0

```
▼ C _____▼  
typedef struct pmix_cpuset {  
    char *source;  
    void *bitmap;  
} pmix_cpuset_t;  
▲ C _____▲
```

9 17.2.18.2 Cpuset support macros

10 The following macros support the `pmix_cpuset_t` structure.

11 **Static initializer for the cpuset structure**

12 *Provisional*

Provide a static initializer for the `pmix_cpuset_t` fields.

PMIx v4.2

```
▼ C _____▼  
PMIX_CPUSSET_STATIC_INIT  
▲ C _____▲
```

14 **Initialize the cpuset structure**

15 Initialize the `pmix_cpuset_t` fields.

PMIx v4.0

```
▼ C _____▼  
PMIX_CPUSSET_CONSTRUCT (m)  
▲ C _____▲
```

17 **IN** m

18 Pointer to the structure to be initialized (pointer to `pmix_cpuset_t`)

19 **Destruct the cpuset structure**

20 Destruct the `pmix_cpuset_t` fields.

PMIx v4.0

```
▼ C _____▼  
PMIX_CPUSSET_DESTRUCT (m)  
▲ C _____▲
```

22 **IN** m

23 Pointer to the structure to be destructed (pointer to `pmix_cpuset_t`)

1 **Create a cpuset array**
2 Allocate and initialize a `pmix_cpuset_t` array.

▼ `PMIX_CPUSSET_CREATE` C

3 **PMIX_CPUSSET_CREATE** (*m*, *n*)

▲ `PMIX_CPUSSET_CREATE` C

4 **INOUT** *m*

5 Address where the pointer to the array of `pmix_cpuset_t` structures shall be stored
6 (handle)

7 **IN** *n*

8 Number of structures to be allocated (size_t)

9 **Release a cpuset array**

10 Deconstruct and free a `pmix_cpuset_t` array.

PMIx v4.0

▼ `PMIX_CPUSSET_FREE` C

11 **PMIX_CPUSSET_FREE** (*m*, *n*)

▲ `PMIX_CPUSSET_FREE` C

12 **INOUT** *m*

13 Address the array of `pmix_cpuset_t` structures to be released (handle)

14 **IN** *n*

15 Number of structures in the array (size_t)

16 17.2.19 `PMIx_server_define_process_set`

17 **Summary**

18 Define a `PMIx` process set.

19 **Format**

PMIx v4.0

▼ `PMIx_server_define_process_set` C

20 `pmix_status_t`

21 `PMIx_server_define_process_set` (const `pmix_proc_t` *members*[],

22 `size_t` *nmembers*,

23 `char *`*pset_name*);

▲ `PMIx_server_define_process_set` C

24 **IN** *members*

25 Pointer to an array of `pmix_proc_t` containing the identifiers of the processes in the
26 process set (handle)

27 **IN** *nmembers*

28 Number of elements in *members* (integer)

29 **IN** *pset_name*

30 String name of the process set being defined (`char*`)

31 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Description

Provide a function by which the host environment can create a process set. The PMIx server shall alert all local clients of the new process set (including process set name and membership) via the [PMIX_PROCESS_SET_DEFINE](#) event.

Advice to PMIx server hosts

The host environment is responsible for ensuring:

- consistent knowledge of process set membership across all involved PMIx servers; and
- that process set names do not conflict with system-assigned namespaces within the scope of the set

17.2.20 PMIx_server_delete_process_set

Summary

Delete a PMIx process set name

Format

PMIx v4.0

```
pmix_status_t  
PMIx_server_delete_process_set(char *pset_name);
```

IN pset_name

String name of the process set being deleted (**char***)

Returns [PMIX_SUCCESS](#) or a negative value indicating the error.

Description

Provide a function by which the host environment can delete a process set name. The PMIx server shall alert all local clients of the process set name being deleted via the [PMIX_PROCESS_SET_DELETE](#) event. Deletion of the name has no impact on the member processes.

Advice to PMIx server hosts

The host environment is responsible for ensuring consistent knowledge of process set membership across all involved PMIx servers.

1 17.3 Server Function Pointers

2 PMIx utilizes a "function-shipping" approach to support for implementing the server-side of the
3 protocol. This method allows RMs to implement the server without being burdened with PMIx
4 internal details. When a request is received from the client, the corresponding server function will
5 be called with the information.

6 Any functions not supported by the RM can be indicated by a **NULL** for the function pointer. PMIx
7 implementations are required to return a **PMIX_ERR_NOT_SUPPORTED** status to all calls to
8 functions that require host environment support and are not backed by a corresponding server
9 module entry. Host environments may, if they choose, include a function pointer for operations they
10 have not yet implemented and simply return **PMIX_ERR_NOT_SUPPORTED**.

11 Functions that accept directives (i.e., arrays of **pmix_info_t** structures) must check any provided
12 directives for those marked as *required* via the **PMIX_INFO_REQD** flag. PMIx client and server
13 libraries are required to mark any such directives with the **PMIX_INFO_REQD_PROCESSED** flag
14 should they have handled the request. Any required directive that has not been marked therefore
15 becomes the responsibility of the host environment. If a required directive that hasn't been
16 processed by a lower level cannot be supported by the host, then the
17 **PMIX_ERR_NOT_SUPPORTED** error constant must be returned. If the directive can be processed
18 by the host, then the host shall do so and mark the attribute with the
19 **PMIX_INFO_REQD_PROCESSED** flag.

20 The host RM will provide the function pointers in a **pmix_server_module_t** structure passed
21 to **PMIx_server_init**. The module structure and associated function references are defined in
22 this section.

▼ Advice to PMIx server hosts ▼

23 For performance purposes, the host server is required to return as quickly as possible from all
24 functions. Execution of the function is thus to be done asynchronously so as to allow the PMIx
25 server support library to handle multiple client requests as quickly and scalably as possible.

26 All data passed to the host server functions is "owned" by the PMIx server support library and
27 must not be free'd. Data returned by the host server via callback function is owned by the host
28 server, which is free to release it upon return from the callback

29 17.3.1 **pmix_server_module_t** Module

30 Summary

31 List of function pointers that a PMIx server passes to **PMIx_server_init** during startup.

Format

C

```
1 typedef struct pmix_server_module_4_0_0_t {
2     /* v1x interfaces */
3     pmix_server_client_connected_fn_t    client_connected; // DEPRECATED
4     pmix_server_client_finalized_fn_t    client_finalized;
5     pmix_server_abort_fn_t               abort;
6     pmix_server_fence_nb_fn_t            fence_nb;
7     pmix_server_dmodex_req_fn_t           direct_modex;
8     pmix_server_publish_fn_t              publish;
9     pmix_server_lookup_fn_t               lookup;
10    pmix_server_unpublish_fn_t             unpublish;
11    pmix_server_spawn_fn_t                 spawn;
12    pmix_server_connect_fn_t               connect;
13    pmix_server_disconnect_fn_t            disconnect;
14    pmix_server_register_events_fn_t       register_events;
15    pmix_server_deregister_events_fn_t     deregister_events;
16    pmix_server_listener_fn_t              listener;
17    /* v2x interfaces */
18    pmix_server_notify_event_fn_t          notify_event;
19    pmix_server_query_fn_t                  query;
20    pmix_server_tool_connection_fn_t       tool_connected;
21    pmix_server_log_fn_t                    log;
22    pmix_server_alloc_fn_t                  allocate;
23    pmix_server_job_control_fn_t            job_control;
24    pmix_server_monitor_fn_t               monitor;
25    /* v3x interfaces */
26    pmix_server_get_cred_fn_t               get_credential;
27    pmix_server_validate_cred_fn_t          validate_credential;
28    pmix_server_iof_fn_t                    iof_pull;
29    pmix_server_stdin_fn_t                  push_stdin;
30    /* v4x interfaces */
31    pmix_server_grp_fn_t                    group;
32    pmix_server_fabric_fn_t                 fabric;
33    pmix_server_client_connected2_fn_t     client_connected2;
34 } pmix_server_module_t;
```

C

Advice to PMIx server hosts

Note that some PMIx implementations *require* the use of C99-style designated initializers to clearly correlate each provided function pointer with the correct member of the `pmix_server_module_t` structure as the location/ordering of struct members may change over time.

1 17.3.2 pmix_server_client_connected_fn_t

2 Summary

3 Notify the host server that a client connected to this server. This function module entry has been
4 **DEPRECATED** in favor of [pmix_server_client_connected2_fn_t](#).

5 Format

PMIx v1.0

C

```
6 typedef pmix_status_t (*pmix_server_client_connected_fn_t) (  
7     const pmix_proc_t *proc,  
8     void* server_object,  
9     pmix_op_cbfunc_t cbfunc,  
10    void *cbdata);
```

C

11 IN proc

12 [pmix_proc_t](#) structure (handle)

13 IN server_object

14 object reference (memory reference)

15 IN cbfunc

16 Callback function [pmix_op_cbfunc_t](#) (function reference)

17 IN cbdata

18 Data to be passed to the callback function (memory reference)

19 Returns one of the following:

- 20 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
21 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
22 prior to returning from the API.
- 23 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
24 returned *success* - the *cbfunc* will not be called
- 25 • a PMIx error constant indicating either an error in the input or that the request was immediately
26 processed and failed - the *cbfunc* will not be called

27 Description

28 This function module entry has been DEPRECATED in favor of
29 [pmix_server_client_connected2_fn_t](#). If both functions are provided, the PMIx
30 library will ignore this function module entry in favor of its replacement.

1 17.3.3 pmix_server_client_connected2_fn_t

2 Summary

3 Notify the host server that a client connected to this server - this version of the original function
4 definition has been extended to include an array of `pmix_info_t`, thereby allowing the PMIx
5 server library to pass additional information identifying the client to the host environment.

6 *PMIx v4.0* Format

```
7 typedef pmix_status_t (*pmix_server_client_connected2_fn_t) (  
8     const pmix_proc_t *proc,  
9     void* server_object,  
10    pmix_info_t info[], size_t ninfo,  
11    pmix_op_cbfunc_t cbfunc,  
12    void *cbdata)  
13
```

13 IN `proc`

14 `pmix_proc_t` structure (handle)

15 IN `server_object`

16 object reference (memory reference)

17 IN `info`

18 Array of info structures (array of handles)

19 IN `ninfo`

20 Number of elements in the *info* array (integer)

21 IN `cbfunc`

22 Callback function `pmix_op_cbfunc_t` (function reference)

23 IN `cbdata`

24 Data to be passed to the callback function (memory reference)

25 Returns one of the following:

- 26 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
27 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
28 prior to returning from the API.
- 29 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
30 returned *success* - the *cbfunc* will not be called
- 31 • a PMIx error constant indicating either an error in the input or that the request was immediately
32 processed and failed - the *cbfunc* will not be called. The PMIx server library is to immediately
33 terminate the connection.

Description

Notify the host environment that a client has called `PMIx_Init`. Note that the client will be in a blocked state until the host server executes the callback function, thus allowing the PMIx server support library to release the client. The `server_object` parameter will be the value of the `server_object` parameter passed to `PMIx_server_register_client` by the host server when registering the connecting client. A host server can choose to not be notified when clients connect by setting `pmix_server_client_connected2_fn_t` to `NULL`.

It is possible that only a subset of the clients in a namespace call `PMIx_Init`. The server's `pmix_server_client_connected2_fn_t` implementation should therefore not depend on being called once per rank in a namespace or delay calling the callback function until all ranks have connected. However, the host may rely on the `pmix_server_client_connected2_fn_t` function module entry being called for a given rank prior to any other function module entries being executed on behalf of that rank.

17.3.4 `pmix_server_client_finalized_fn_t`

Summary

Notify the host environment that a client called `PMIx_Finalize`.

Format

PMIx v1.0

```
typedef pmix_status_t (*pmix_server_client_finalized_fn_t) (  
    const pmix_proc_t *proc,  
    void* server_object,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata);
```

IN proc
`pmix_proc_t` structure (handle)

IN server_object
object reference (memory reference)

IN cbfunc
Callback function `pmix_op_cbfunc_t` (function reference)

IN cbdata
Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned `success` - the `cbfunc` will not be called

- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Description

Notify the host environment that a client called `PMIx_Finalize`. Note that the client will be in a blocked state until the host server executes the callback function, thus allowing the PMIx server support library to release the client. The `server_object` parameter will be the value of the `server_object` parameter passed to `PMIx_server_register_client` by the host server when registering the connecting client. If provided, an implementation of `pmix_server_client_finalized_fn_t` is only required to call the callback function designated. A host server can choose to not be notified when clients finalize by setting `pmix_server_client_finalized_fn_t` to `NULL`.

Note that the host server is only being informed that the client has called `PMIx_Finalize`. The client might not have exited. If a client exits without calling `PMIx_Finalize`, the server support library will not call the `pmix_server_client_finalized_fn_t` implementation.

Advice to PMIx server hosts

This operation is an opportunity for a host server to update the status of the tasks it manages. It is also a convenient and well defined time to release resources used to support that client.

17.3.5 pmix_server_abort_fn_t

Summary

Notify the host environment that a local client called `PMIx_Abort`.

Format

C

```
typedef pmix_status_t (*pmix_server_abort_fn_t) (
    const pmix_proc_t *proc,
    void *server_object,
    int status,
    const char msg[],
    pmix_proc_t procs[],
    size_t nprocs,
    pmix_op_cbfunc_t cbfunc,
    void *cbdata);
```


1 **IN** **proc**
2 **pmix_proc_t** structure identifying the process requesting the abort (handle)
3 **IN** **server_object**
4 object reference (memory reference)
5 **IN** **status**
6 exit status (integer)
7 **IN** **msg**
8 exit status message (string)
9 **IN** **procs**
10 Array of **pmix_proc_t** structures identifying the processes to be terminated (array of
11 handles)
12 **IN** **nprocs**
13 Number of elements in the *procs* array (integer)
14 **IN** **cbfunc**
15 Callback function **pmix_op_cbfunc_t** (function reference)
16 **IN** **cbdata**
17 Data to be passed to the callback function (memory reference)

18 Returns one of the following:

- 19 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
20 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
21 prior to returning from the API.
- 22 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
23 returned *success* - the *cbfunc* will not be called
- 24 • **PMIX_ERR_PARAM_VALUE_NOT_SUPPORTED** indicating that the host environment supports
25 this API, but the request includes processes that the host environment cannot abort - e.g., if the
26 request is to abort subsets of processes from a namespace, or processes outside of the caller's
27 own namespace, and the host environment does not permit such operations. In this case, none of
28 the specified processes will be terminated - the *cbfunc* will not be called
- 29 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
30 request, even though the function entry was provided in the server module - the *cbfunc* will not
31 be called
- 32 • a PMIx error constant indicating either an error in the input or that the request was immediately
33 processed and failed - the *cbfunc* will not be called

Description

A local client called `PMIx_Abort`. Note that the client will be in a blocked state until the host server executes the callback function, thus allowing the PMIx server library to release the client. The array of `procs` indicates which processes are to be terminated. A `NULL` for the `procs` array indicates that all processes in the caller's namespace are to be aborted, including itself - this is the equivalent of passing a `pmix_proc_t` array element containing the caller's namespace and a rank value of `PMIX_RANK_WILDCARD`.

17.3.6 pmix_server_fencebn_fn_t

Summary

At least one client called either `PMIx_Fence` or `PMIx_Fence_nb`.

Format

PMIx v1.0

C

```
typedef pmix_status_t (*pmix_server_fencebn_fn_t) (  
    const pmix_proc_t procs[],  
    size_t nprocs,  
    const pmix_info_t info[],  
    size_t ninfo,  
    char *data, size_t ndata,  
    pmix_modex_cbfunc_t cbfunc,  
    void *cbdata);
```

C

- IN procs**
Array of `pmix_proc_t` structures identifying operation participants(array of handles)
- IN nprocs**
Number of elements in the `procs` array (integer)
- IN info**
Array of info structures (array of handles)
- IN ninfo**
Number of elements in the `info` array (integer)
- IN data**
(string)
- IN ndata**
(integer)
- IN cbfunc**
Callback function `pmix_modex_cbfunc_t` (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

Returns one of the following:

- 1 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
2 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
3 prior to returning from the API.
- 4 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
5 request, even though the function entry was provided in the server module - the *cbfunc* will not
6 be called
- 7 • a PMIx error constant indicating either an error in the input or that the request was immediately
8 processed and failed - the *cbfunc* will not be called

Required Attributes

9 PMIx libraries are required to pass any provided attributes to the host environment for processing.

10 The following attributes are required to be supported by all host environments:

11 **PMIX_COLLECT_DATA** "**pmix.collect**" (**bool**)

12 Collect all data posted by the participants using **PMIx_Put** that has been committed via
13 **PMIx_Commit**, making the collection locally available to each participant at the end of the
14 operation. By default, this will include all job-level information that was locally generated
15 by PMIx servers unless excluded using the **PMIX_COLLECT_GENERATED_JOB_INFO**
16 attribute.

17 **PMIX_LOCAL_COLLECTIVE_STATUS** "**pmix.loc.col.st**" (**pmix_status_t**)

18 Status code for local collective operation being reported to the host by the server library.
19 PMIx servers may aggregate the participation by local client processes in a collective
20 operation - e.g., instead of passing individual client calls to **PMIx_Fence** up to the host
21 environment, the server may pass only a single call to the host when all local participants
22 have executed their **PMIx_Fence** call, thereby reducing the burden placed on the host.
23 However, in cases where the operation locally fails (e.g., if a participating client abnormally
24 terminates prior to calling the operation), the server upcall functions to the host do not
25 include a **pmix_status_t** by which the PMIx server can alert the host to that failure.
26 This attribute resolves that problem by allowing the server to pass the status information
27 regarding the local collective operation.

Optional Attributes

28 The following attributes are optional for host environments:

29 **PMIX_TIMEOUT** "**pmix.timeout**" (**int**)

30 Time in seconds before the specified operation should time out (zero indicating infinite) and
31 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
32 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Advice to PMIx server hosts

1 Host environment are required to return `PMIX_ERR_NOT_SUPPORTED` if passed an attributed
2 marked as `PMIX_INFO_REQD` that they do not support, even if support for that attribute is
3 optional.

Description

4 All local clients in the provided array of *procs* called either `PMIx_Fence` or `PMIx_Fence_nb`.
5 In either case, the host server will be called via a non-blocking function to execute the specified
6 operation once all participating local processes have contributed. All processes in the specified
7 *procs* array are required to participate in the `PMIx_Fence/PMIx_Fence_nb` operation. The
8 callback is to be executed once every daemon hosting at least one participant has called the host
9 server's `pmix_server_fencefn_t` function.

11 The provided data is to be collectively shared with all PMIx servers involved in the fence operation,
12 and returned in the modex *cbfunc*. A `NULL` data value indicates that the local processes had no data
13 to contribute.

14 The array of *info* structs is used to pass user-requested options to the server. This can include
15 directives as to the algorithm to be used to execute the fence operation. The directives are optional
16 unless the `PMIX_INFO_REQD` flag has been set - in such cases, the host RM is required to return
17 an error if the directive cannot be met.

Advice to PMIx library implementers

18 The PMIx server library is required to aggregate participation by local clients, passing the request
19 to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

20 The host will receive a single call for each collective operation. It is the responsibility of the host to
21 identify the nodes containing participating processes, execute the collective across all participating
22 nodes, and notify the local PMIx server library upon completion of the global collective. Data
23 received from each node must be simply concatenated to form an aggregated unit, as shown in the
24 following example:

```
25 uint8_t *blob1, *blob2, *total;  
26 size_t sz_blob1, sz_blob2, sz_total;  
27  
28 sz_total = sz_blob1 + sz_blob2;  
29 total = (uint8_t*)malloc(sz_total);  
30 memcpy(total, blob1, sz_blob1);  
31 memcpy(&total[sz_blob1], blob2, sz_blob2);
```

C

1 Note that the ordering of the data blobs does not matter. The host is responsible for free'ing the
2 *data* object passed to it by the PMIx server library.

3 17.3.6.1 Modex Callback Function

4 Summary

5 The `pmix_modex_cbfunc_t` is used by the `pmix_server_fence_fn_t` and
6 `pmix_server_dmodex_req_fn_t` PMIx server operations to return modex Business Card
7 Exchange (BCX) data.

PMIx v1.0

C

```
8 typedef void (*pmix_modex_cbfunc_t)
9     (pmix_status_t status,
10      const char *data, size_t ndata,
11      void *cbdata,
12      pmix_release_cbfunc_t release_fn,
13      void *release_cbdata);
```

C

14 **IN status**
15 Status associated with the operation (handle)

16 **IN data**
17 Data to be passed (pointer)

18 **IN ndata**
19 size of the data (`size_t`)

20 **IN cbdata**
21 Callback data passed to original API call (memory reference)

22 **IN release_fn**
23 Callback for releasing *data* (function pointer)

24 **IN release_cbdata**
25 Pointer to be passed to *release_fn* (memory reference)

26 Description

27 A callback function that is solely used by PMIx servers, and not clients, to return modex BCX data
28 in response to “fence” and “get” operations. The returned blob contains the data collected from
29 each server participating in the operation.

30 17.3.7 `pmix_server_dmodex_req_fn_t`

31 Summary

32 Used by the PMIx server to request its local host contact the PMIx server on the remote node that
33 hosts the specified process to obtain and return a direct modex blob for that process.

Format

C

```
typedef pmix_status_t (*pmix_server_dmodex_req_fn_t) (  
    const pmix_proc_t *proc,  
    const pmix_info_t info[],  
    size_t ninfo,  
    pmix_modex_cbfunc_t cbfunc,  
    void *cbdata);
```

C

- IN** `proc`
`pmix_proc_t` structure identifying the process whose data is being requested (handle)
- IN** `info`
Array of info structures (array of handles)
- IN** `ninfo`
Number of elements in the `info` array (integer)
- IN** `cbfunc`
Callback function `pmix_modex_cbfunc_t` (function reference)
- IN** `cbdata`
Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function prior to returning from the API.
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the `cbfunc` will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the `cbfunc` will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing.

All host environments are required to support the following attributes:

- PMIX_REQUIRED_KEY** "pmix.req.key" (`char*`)
Identifies a key that must be included in the requested information. If the specified key is not already available, then the PMIx servers are required to delay response to the `dmodex` request until either the key becomes available or the request times out.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Used by the PMIx server to request its local host contact the PMIx server on the remote node that hosts the specified proc to obtain and return any information that process posted via calls to **PMIx_Put** and **PMIx_Commit**.

The array of *info* structs is used to pass user-requested options to the server. This can include a timeout to preclude an indefinite wait for data that may never become available. The directives are optional unless the *mandatory* flag has been set - in such cases, the host RM is required to return an error if the directive cannot be met.

17.3.7.1 Dmodex attributes

PMIX_REQUIRED_KEY "pmix.req.key" (char*)

Identifies a key that must be included in the requested information. If the specified key is not already available, then the PMIx servers are required to delay response to the dmodex request until either the key becomes available or the request times out.

17.3.8 pmix_server_publish_fn_t

Summary

Publish data per the PMIx API specification.

Format

C

```
typedef pmix_status_t (*pmix_server_publish_fn_t) (
    const pmix_proc_t *proc,
    const pmix_info_t info[],
    size_t ninfo,
    pmix_op_cbfunc_t cbfunc,
    void *cbdata);
```

1 **IN** `proc`
 2 `pmix_proc_t` structure of the process publishing the data (handle)
 3 **IN** `info`
 4 Array of info structures (array of handles)
 5 **IN** `ninfo`
 6 Number of elements in the `info` array (integer)
 7 **IN** `cbfunc`
 8 Callback function `pmix_op_cbfunc_t` (function reference)
 9 **IN** `cbdata`
 10 Data to be passed to the callback function (memory reference)

11 Returns one of the following:

- 12 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
 13 will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function
 14 prior to returning from the API.
- 15 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
 16 returned `success` - the `cbfunc` will not be called
- 17 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
 18 request, even though the function entry was provided in the server module - the `cbfunc` will not
 19 be called
- 20 • a PMIx error constant indicating either an error in the input or that the request was immediately
 21 processed and failed - the `cbfunc` will not be called

Required Attributes

22 PMIx libraries are required to pass any provided attributes to the host environment for processing.
 23 In addition, the following attributes are required to be included in the passed `info` array:

24 **PMIX_USERID** "`pmix.euid`" (`uint32_t`)

25 Effective user ID of the connecting process.

26 **PMIX_GRPID** "`pmix.egid`" (`uint32_t`)

27 Effective group ID of the connecting process.

28
 29 Host environments that implement this entry point are required to support the following attributes:

30 **PMIX_RANGE** "`pmix.range`" (`pmix_data_range_t`)

31 Define constraints on the processes that can access published data or generated events or
 32 define constraints on the provider of data when looking up published data.

33 **PMIX_PERSISTENCE** "`pmix.persist`" (`pmix_persistence_t`)

1 Declare how long the datastore shall retain the provided data. The datastore is to delete the
2 data upon reaching the persistence criterion.

Optional Attributes

3 The following attributes are optional for host environments that support this operation:

4 **PMIX_TIMEOUT** "pmix.timeout" (int)

5 Time in seconds before the specified operation should time out (zero indicating infinite) and
6 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
7 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

8 Publish data per the **PMIx_Publish** specification. The callback is to be executed upon
9 completion of the operation. The default data range is left to the host environment, but expected to
10 be **PMIX_RANGE_SESSION**, and the default persistence **PMIX_PERSIST_SESSION** or their
11 equivalent. These values can be specified by including the respective attributed in the *info* array.

12 The persistence indicates how long the server should retain the data.

Advice to PMIx server hosts

13 The host environment is not required to guarantee support for any specific range - i.e., the
14 environment does not need to return an error if the data store doesn't support a specified range so
15 long as it is covered by some internally defined range. However, the server must return an error (a)
16 if the key is duplicative within the storage range, and (b) if the server does not allow overwriting of
17 published info by the original publisher - it is left to the discretion of the host environment to allow
18 info-key-based flags to modify this behavior.

19 The **PMIX_USERID** and **PMIX_GRPID** of the publishing process will be provided to support
20 authorization-based access to published information and must be returned on any subsequent
21 lookup request.
22

23 17.3.9 pmix_server_lookup_fn_t

24 Summary

25 Lookup published data.

Format

```
typedef pmix_status_t (*pmix_server_lookup_fn_t) (
    const pmix_proc_t *proc,
    char **keys,
    const pmix_info_t info[],
    size_t ninfo,
    pmix_lookup_cbfunc_t cbfunc,
    void *cbdata);
```

IN `proc`

`pmix_proc_t` structure of the process seeking the data (handle)

IN `keys`

(array of strings)

IN `info`

Array of info structures (array of handles)

IN `ninfo`

Number of elements in the *info* array (integer)

IN `cbfunc`

Callback function `pmix_lookup_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (`uint32_t`)

Effective user ID of the connecting process.

1 **PMIX_GRPID** "pmix.egid" (uint32_t)
2 Effective group ID of the connecting process.

3
4 Host environments that implement this entry point are required to support the following attributes:

5 **PMIX_RANGE** "pmix.range" (pmix_data_range_t)
6 Define constraints on the processes that can access published data or generated events or
7 define constraints on the provider of data when looking up published data.

8 **PMIX_WAIT** "pmix.wait" (int)
9 Caller requests that the PMIx server wait until at least the specified number of values are
10 found (a value of zero indicates *all* and is the default).

11
12  **Optional Attributes**

11 The following attributes are optional for host environments that support this operation:

12 **PMIX_TIMEOUT** "pmix.timeout" (int)
13 Time in seconds before the specified operation should time out (zero indicating infinite) and
14 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
15 caused by multiple layers (client, server, and host) simultaneously timing the operation.

16 **Description**

17 Lookup published data. The host server will be passed a **NULL**-terminated array of string keys
18 identifying the data being requested.

19 The array of *info* structs is used to pass user-requested options to the server. The default data range
20 is left to the host environment, but expected to be **PMIX_RANGE_SESSION**. This can include a
21 wait flag to indicate that the server should wait for all data to become available before executing the
22 callback function, or should immediately callback with whatever data is available. In addition, a
23 timeout can be specified on the wait to preclude an indefinite wait for data that may never be
24 published.

25  **Advice to PMIx server hosts**

25 The **PMIX_USERID** and **PMIX_GRPID** of the requesting process will be provided to support
26 authorization-based access to published information. The host environment is not required to
27 guarantee support for any specific range - i.e., the environment does not need to return an error if
28 the data store doesn't support a specified range so long as it is covered by some internally defined
29 range.

1 17.3.10 pmix_server_unpublish_fn_t

2 Summary

3 Delete data from the data store.

4 Format

PMIx v1.0

C

```
5 typedef pmix_status_t (*pmix_server_unpublish_fn_t) (  
6     const pmix_proc_t *proc,  
7     char **keys,  
8     const pmix_info_t info[],  
9     size_t ninfo,  
10    pmix_op_cbfunc_t cbfunc,  
11    void *cbdata);
```

C

12 **IN** `proc`
13 `pmix_proc_t` structure identifying the process making the request (handle)

14 **IN** `keys`
15 (array of strings)

16 **IN** `info`
17 Array of info structures (array of handles)

18 **IN** `ninfo`
19 Number of elements in the *info* array (integer)

20 **IN** `cbfunc`
21 Callback function `pmix_op_cbfunc_t` (function reference)

22 **IN** `cbdata`
23 Data to be passed to the callback function (memory reference)

24 Returns one of the following:

- 25 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
26 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
27 prior to returning from the API.
- 28 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
29 returned *success* - the *cbfunc* will not be called
- 30 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
31 request, even though the function entry was provided in the server module - the *cbfunc* will not
32 be called
- 33 • a PMIx error constant indicating either an error in the input or that the request was immediately
34 processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (uint32_t)

Effective user ID of the connecting process.

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group ID of the connecting process.

Host environments that implement this entry point are required to support the following attributes:

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Define constraints on the processes that can access published data or generated events or define constraints on the provider of data when looking up published data.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Delete data from the data store. The host server will be passed a **NULL**-terminated array of string keys, plus potential directives such as the data range within which the keys should be deleted. The default data range is left to the host environment, but expected to be **PMIX_RANGE_SESSION**. The callback is to be executed upon completion of the delete procedure.

Advice to PMIx server hosts

The **PMIX_USERID** and **PMIX_GRPID** of the requesting process will be provided to support authorization-based access to published information. The host environment is not required to guarantee support for any specific range - i.e., the environment does not need to return an error if the data store doesn't support a specified range so long as it is covered by some internally defined range.

1 17.3.11 pmix_server_spawn_fn_t

2 Summary

3 Spawn a set of applications/processes as per the [PMIx_Spawn](#) API.

4 Format

C

```
5 typedef pmix_status_t (*pmix_server_spawn_fn_t) (  
6     const pmix_proc_t *proc,  
7     const pmix_info_t job_info[],  
8     size_t ninfo,  
9     const pmix_app_t apps[],  
10    size_t napps,  
11    pmix_spawn_cbfunc_t cbfunc,  
12    void *cbdata);
```

C

- 13 **IN** `proc`
14 [pmix_proc_t](#) structure of the process making the request (handle)
- 15 **IN** `job_info`
16 Array of info structures (array of handles)
- 17 **IN** `ninfo`
18 Number of elements in the *jobinfo* array (integer)
- 19 **IN** `apps`
20 Array of [pmix_app_t](#) structures (array of handles)
- 21 **IN** `napps`
22 Number of elements in the *apps* array (integer)
- 23 **IN** `cbfunc`
24 Callback function [pmix_spawn_cbfunc_t](#) (function reference)
- 25 **IN** `cbdata`
26 Data to be passed to the callback function (memory reference)

27 Returns one of the following:

- 28
- 29 • [PMIX_SUCCESS](#), indicating that the request is being processed by the host environment - result
30 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
31 prior to returning from the API.
 - 32 • [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and
33 returned *success* - the *cbfunc* will not be called
 - 34 • [PMIX_ERR_NOT_SUPPORTED](#), indicating that the host environment does not support the
35 request, even though the function entry was provided in the server module - the *cbfunc* will not
36 be called
 - 37 • a PMIx error constant indicating either an error in the input or that the request was immediately
38 processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx server libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (uint32_t)

Effective user ID of the connecting process.

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group ID of the connecting process.

PMIX_SPAWNED "pmix.spawned" (bool)

true if this process resulted from a call to **PMIx_Spawn**. Lack of inclusion (i.e., a return status of **PMIX_ERR_NOT_FOUND**) corresponds to a value of **false** for this attribute.

PMIX_PARENT_ID "pmix.parent" (pmix_proc_t)

Process identifier of the parent process of the specified process - typically used to identify the application process that caused the job containing the specified process to be spawned (e.g., the process that called **PMIx_Spawn**). This attribute is only provided for a process if it was created by a call to **PMIx_Spawn** or **PMIx_Spawn_nb**.

PMIX_REQUESTOR_IS_TOOL "pmix.req.tool" (bool)

The requesting process is a PMIx tool.

PMIX_REQUESTOR_IS_CLIENT "pmix.req.client" (bool)

The requesting process is a PMIx client.

Host environments that provide this module entry point are required to pass the **PMIX_SPAWNED** and **PMIX_PARENT_ID** attributes to all PMIx servers launching new child processes so those values can be returned to clients upon connection to the PMIx server. In addition, they are required to support the following attributes when present in either the *job_info* or the *info* array of an element of the *apps* array:

PMIX_WDIR "pmix.wdir" (char*)

Working directory for spawned processes.

PMIX_SET_SESSION_CWD "pmix.ssn cwd" (bool)

Set the current working directory to the session working directory assigned by the RM - can be assigned to the entire job (by including attribute in the *job_info* array) or on a per-application basis in the *info* array for each **pmix_app_t**.

PMIX_PREFIX "pmix.prefix" (char*)

Prefix to use for starting spawned processes - i.e., the directory where the executables can be found.

PMIX_HOST "pmix.host" (char*)

Comma-delimited list of hosts to use for spawned processes.

1 **PMIX_HOSTFILE** "pmix.hostfile" (char*)

2 Hostfile to use for spawned processes.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

3 The following attributes are optional for host environments that support this operation:

4 **PMIX_ADD_HOSTFILE** "pmix.addhostfile" (char*)

5 Hostfile containing hosts to add to existing allocation.

6 **PMIX_ADD_HOST** "pmix.addhost" (char*)

7 Comma-delimited list of hosts to add to the allocation.

8 **PMIX_PRELOAD_BIN** "pmix.preloadbin" (bool)

9 Preload executables onto nodes prior to executing launch procedure.

10 **PMIX_PRELOAD_FILES** "pmix.preloadfiles" (char*)

11 Comma-delimited list of files to pre-position on nodes prior to executing launch procedure.

12 **PMIX_PERSONALITY** "pmix.pers" (char*)

13 Name of personality corresponding to programming model used by application - supported
14 values depend upon PMIx implementation.

15 **PMIX_DISPLAY_MAP** "pmix.dispmap" (bool)

16 Display process mapping upon spawn.

17 **PMIX_PPR** "pmix.ppr" (char*)

18 Number of processes to spawn on each identified resource.

19 **PMIX_MAPBY** "pmix.mapby" (char*)

20 Process mapping policy - when accessed using **PMIx_Get**, use the
21 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
22 provided namespace. Supported values are launcher specific.

23 **PMIX_RANKBY** "pmix.rankby" (char*)

24 Process ranking policy - when accessed using **PMIx_Get**, use the
25 **PMIX_RANK_WILDCARD** value for the rank to discover the ranking algorithm used for the
26 provided namespace. Supported values are launcher specific.

27 **PMIX_BINDTO** "pmix.bindto" (char*)

28 Process binding policy - when accessed using **PMIx_Get**, use the
29 **PMIX_RANK_WILDCARD** value for the rank to discover the binding policy used for the
30 provided namespace. Supported values are launcher specific.

31 **PMIX_STDIN_TGT** "pmix.stdin" (uint32_t)

32 Spawned process rank that is to receive any forwarded **stdin**.

33 **PMIX_FWD_STDIN** "pmix.fwd.stdin" (pmix_rank_t)

1 The requester intends to push information from its **stdin** to the indicated process. The
2 local spawn agent should, therefore, ensure that the **stdin** channel to that process remains
3 available. A rank of **PMIX_RANK_WILDCARD** indicates that all processes in the spawned
4 job are potential recipients. The requester will issue a call to **PMIx_IOF_push** to initiate
5 the actual forwarding of information to specified targets - this attribute simply requests that
6 the IL retain the ability to forward the information to the designated targets.

7 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)

8 Requests that the ability to forward the **stdout** of the spawned processes be maintained.
9 The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and
10 other options for delivery of the forwarded output.

11 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)

12 Requests that the ability to forward the **stderr** of the spawned processes be maintained.
13 The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and
14 other options for delivery of the forwarded output.

15 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)

16 Included in the **pmix_info_t** array of a **pmix_app_t**, this attribute declares that the
17 application consists of debugger daemons and shall be governed accordingly. If used as the
18 sole **pmix_app_t** in a **PMIx_Spawn** request, then the **PMIX_DEBUG_TARGET** attribute
19 must also be provided (in either the **job_info** or in the **info** array of the **pmix_app_t**) to
20 identify the namespace to be debugged so that the launcher can determine where to place the
21 spawned daemons. If neither **PMIX_DEBUG_DAEMONS_PER_PROC** nor
22 **PMIX_DEBUG_DAEMONS_PER_NODE** is specified, then the launcher shall default to a
23 placement policy of one daemon per process in the target job.

24 **PMIX_TAG_OUTPUT** "pmix.tagout" (bool)

25 Tag **stdout/stderr** with the identity of the source process - can be assigned to the entire
26 job (by including attribute in the **job_info** array) or on a per-application basis in the **info**
27 array for each **pmix_app_t**.

28 **PMIX_TIMESTAMP_OUTPUT** "pmix.tsout" (bool)

29 Timestamp output - can be assigned to the entire job (by including attribute in the **job_info**
30 array) or on a per-application basis in the **info** array for each **pmix_app_t**.

31 **PMIX_MERGE_STDERR_STDOUT** "pmix.mergeerrout" (bool)

32 Merge **stdout** and **stderr** streams - can be assigned to the entire job (by including
33 attribute in the **job_info** array) or on a per-application basis in the **info** array for each
34 **pmix_app_t**.

35 **PMIX_OUTPUT_TO_FILE** "pmix.outfile" (char*)

36 Direct output (both **stdout** and **stderr**) into files of form "<filename>.rank" - can be
37 assigned to the entire job (by including attribute in the **job_info** array) or on a per-application
38 basis in the **info** array for each **pmix_app_t**.

39 **PMIX_INDEX_ARGV** "pmix.indxargv" (bool)

1 Mark the `argv` with the rank of the process.

2 **PMIX_CPUS_PER_PROC** "`pmix.cpusperproc`" (`uint32_t`)

3 Number of PUs to assign to each rank - when accessed using `PMIx_Get`, use the
4 **PMIX_RANK_WILDCARD** value for the rank to discover the PUs/process assigned to the
5 provided namespace.

6 **PMIX_NO_PROCS_ON_HEAD** "`pmix.nolocal`" (`bool`)

7 Do not place processes on the head node.

8 **PMIX_NO_OVERSUBSCRIBE** "`pmix.noover`" (`bool`)

9 Do not oversubscribe the nodes - i.e., do not place more processes than allocated slots on a
10 node.

11 **PMIX_REPORT_BINDINGS** "`pmix.rebind`" (`bool`)

12 Report bindings of the individual processes.

13 **PMIX_CPU_LIST** "`pmix.cpulist`" (`char*`)

14 List of PUs to use for this job - when accessed using `PMIx_Get`, use the
15 **PMIX_RANK_WILDCARD** value for the rank to discover the PU list used for the provided
16 namespace.

17 **PMIX_JOB_RECOVERABLE** "`pmix.recover`" (`bool`)

18 Application supports recoverable operations.

19 **PMIX_JOB_CONTINUOUS** "`pmix.continuous`" (`bool`)

20 Application is continuous, all failed processes should be immediately restarted.

21 **PMIX_MAX_RESTARTS** "`pmix.maxrestarts`" (`uint32_t`)

22 Maximum number of times to restart a process - when accessed using `PMIx_Get`, use the
23 **PMIX_RANK_WILDCARD** value for the rank to discover the max restarts for the provided
24 namespace.

25 **PMIX_TIMEOUT** "`pmix.timeout`" (`int`)

26 Time in seconds before the specified operation should time out (zero indicating infinite) and
27 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
28 caused by multiple layers (client, server, and host) simultaneously timing the operation.

29 **PMIX_JOB_TIMEOUT** "`pmix.job.time`" (`int`)

30 Time in seconds before the spawned job should time out and be terminated (0 => infinite),
31 defined as the total runtime of the job (equivalent to the walltime limit of typical batch
32 schedulers).

33 **PMIX_SPAWN_TIMEOUT** "`pmix.sp.time`" (`int`)

34 Time in seconds before spawn operation should time out (0 => infinite). Logically
35 equivalent to passing the **PMIX_TIMEOUT** attribute to the `PMIx_Spawn` API, it is
36 provided as a separate attribute to distinguish it from the **PMIX_JOB_TIMEOUT** attribute

Description

Spawn a set of applications/processes as per the [PMIx_Spawn](#) API. Note that applications are not required to be MPI or any other programming model. Thus, the host server cannot make any assumptions as to their required support. The callback function is to be executed once all processes have been started. An error in starting any application or process in this request shall cause all applications and processes in the request to be terminated, and an error returned to the originating caller.

Note that a timeout can be specified in the `job_info` array to indicate that failure to start the requested job within the given time should result in termination to avoid hangs.

17.3.11.1 Server spawn attributes

PMIX_REQUESTOR_IS_TOOL "pmix.req.tool" (bool)

The requesting process is a PMIx tool.

PMIX_REQUESTOR_IS_CLIENT "pmix.req.client" (bool)

The requesting process is a PMIx client.

17.3.12 pmix_server_connect_fn_t

Summary

Record the specified processes as *connected*.

Format

PMIx v1.0

C

```
typedef pmix_status_t (*pmix_server_connect_fn_t) (  
    const pmix_proc_t procs[],  
    size_t nprocs,  
    const pmix_info_t info[],  
    size_t ninfo,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN procs

Array of [pmix_proc_t](#) structures identifying participants (array of handles)

IN nprocs

Number of elements in the *procs* array (integer)

IN info

Array of info structures (array of handles)

IN ninfo

Number of elements in the *info* array (integer)

IN cbfunc

Callback function [pmix_op_cbfunc_t](#) (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

1 Returns one of the following:

- 2 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
3 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
4 prior to returning from the API.
- 5 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
6 returned *success* - the *cbfunc* will not be called
- 7 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
8 request, even though the function entry was provided in the server module - the *cbfunc* will not
9 be called
- 10 • a PMIx error constant indicating either an error in the input or that the request was immediately
11 processed and failed - the *cbfunc* will not be called

Required Attributes

12 **PMIX_LOCAL_COLLECTIVE_STATUS** "`pmix.loc.col.st`" (`pmix_status_t`)

13 Status code for local collective operation being reported to the host by the server library.
14 PMIx servers may aggregate the participation by local client processes in a collective
15 operation - e.g., instead of passing individual client calls to **PMIx_Fence** up to the host
16 environment, the server may pass only a single call to the host when all local participants
17 have executed their **PMIx_Fence** call, thereby reducing the burden placed on the host.
18 However, in cases where the operation locally fails (e.g., if a participating client abnormally
19 terminates prior to calling the operation), the server upcall functions to the host do not
20 include a `pmix_status_t` by which the PMIx server can alert the host to that failure.
21 This attribute resolves that problem by allowing the server to pass the status information
22 regarding the local collective operation.

23 PMIx libraries are required to pass any provided attributes to the host environment for processing.

Optional Attributes

24 The following attributes are optional for host environments that support this operation:

25 **PMIX_TIMEOUT** "`pmix.timeout`" (`int`)

26 Time in seconds before the specified operation should time out (zero indicating infinite) and
27 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
28 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Record the processes specified by the *procs* array as *connected* as per the PMIx definition. The callback is to be executed once every daemon hosting at least one participant has called the host server's `pmix_server_connect_fn_t` function, and the host environment has completed any supporting operations required to meet the terms of the PMIx definition of *connected* processes.

Advice to PMIx library implementers

The PMIx server library is required to aggregate participation by local clients, passing the request to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

The host will receive a single call for each collective operation. It is the responsibility of the host to identify the nodes containing participating processes, execute the collective across all participating nodes, and notify the local PMIx server library upon completion of the global collective.

17.3.13 `pmix_server_disconnect_fn_t`

Summary

Disconnect a previously connected set of processes.

Format

C

```
typedef pmix_status_t (*pmix_server_disconnect_fn_t) (  
    const pmix_proc_t procs[],  
    size_t nprocs,  
    const pmix_info_t info[],  
    size_t ninfo,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN `procs`

Array of `pmix_proc_t` structures identifying participants (array of handles)

IN `nprocs`

Number of elements in the `procs` array (integer)

IN `info`

Array of info structures (array of handles)

IN `ninfo`

Number of elements in the `info` array (integer)

IN `cbfunc`

Callback function `pmix_op_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the `cbfunc` will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the `cbfunc` will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the `cbfunc` will not be called

Required Attributes

PMIX_LOCAL_COLLECTIVE_STATUS "pmix.loc.col.st" (`pmix_status_t`)

1 Status code for local collective operation being reported to the host by the server library.
2 PMIx servers may aggregate the participation by local client processes in a collective
3 operation - e.g., instead of passing individual client calls to **PMIx_Fence** up to the host
4 environment, the server may pass only a single call to the host when all local participants
5 have executed their **PMIx_Fence** call, thereby reducing the burden placed on the host.
6 However, in cases where the operation locally fails (e.g., if a participating client abnormally
7 terminates prior to calling the operation), the server upcall functions to the host do not
8 include a **pmix_status_t** by which the PMIx server can alert the host to that failure.
9 This attribute resolves that problem by allowing the server to pass the status information
10 regarding the local collective operation.

11 PMIx libraries are required to pass any provided attributes to the host environment for processing.

Optional Attributes

12 The following attributes are optional for host environments that support this operation:

13 **PMIX_TIMEOUT** "pmix.timeout" (int)

14 Time in seconds before the specified operation should time out (zero indicating infinite) and
15 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
16 caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

17 Disconnect a previously connected set of processes. The callback is to be executed once every
18 daemon hosting at least one participant has called the host server's has called the
19 **pmix_server_disconnect_fn_t** function, and the host environment has completed any
20 required supporting operations.
21

Advice to PMIx library implementers

22 The PMIx server library is required to aggregate participation by local clients, passing the request
23 to the host environment once all local participants have executed the API.

Advice to PMIx server hosts

24 The host will receive a single call for each collective operation. It is the responsibility of the host to
25 identify the nodes containing participating processes, execute the collective across all participating
26 nodes, and notify the local PMIx server library upon completion of the global collective.

27 A **PMIX_ERR_INVALID_OPERATION** error must be returned if the specified set of *procs* was
28 not previously *connected* via a call to the **pmix_server_connect_fn_t** function.

1 17.3.14 pmix_server_register_events_fn_t

2 Summary

3 Register to receive notifications for the specified events.

4 Format

PMIx v1.0

C

```
5 typedef pmix_status_t (*pmix_server_register_events_fn_t) (  
6     pmix_status_t *codes,  
7     size_t ncodes,  
8     const pmix_info_t info[],  
9     size_t ninfo,  
10    pmix_op_cbfunc_t cbfunc,  
11    void *cbdata);
```

C

- 12 **IN codes**
13 Array of [pmix_status_t](#) values (array of handles)
- 14 **IN ncodes**
15 Number of elements in the *codes* array (integer)
- 16 **IN info**
17 Array of info structures (array of handles)
- 18 **IN ninfo**
19 Number of elements in the *info* array (integer)
- 20 **IN cbfunc**
21 Callback function [pmix_op_cbfunc_t](#) (function reference)
- 22 **IN cbdata**
23 Data to be passed to the callback function (memory reference)

24 Returns one of the following:

- 25 • [PMIX_SUCCESS](#), indicating that the request is being processed by the host environment - result
26 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
27 prior to returning from the API.
- 28 • [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and
29 returned *success* - the *cbfunc* will not be called
- 30 • [PMIX_ERR_NOT_SUPPORTED](#), indicating that the host environment does not support the
31 request, even though the function entry was provided in the server module - the *cbfunc* will not
32 be called
- 33 • a PMIx error constant indicating either an error in the input or that the request was immediately
34 processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (uint32_t)

Effective user ID of the connecting process.

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group ID of the connecting process.

Description

Register to receive notifications for the specified status codes. The *info* array included in this API is reserved for possible future directives to further steer notification.

Advice to PMIx library implementers

The PMIx server library must track all client registrations for subsequent notification. This module function shall only be called when:

- the client has requested notification of an environmental code (i.e., a PMIx codes in the range between **PMIX_EVENT_SYS_BASE** and **PMIX_EVENT_SYS_OTHER**, inclusive) or codes that lies outside the defined PMIx range of constants; and
- the PMIx server library has not previously requested notification of that code - i.e., the host environment is to be contacted only once a given unique code value

Advice to PMIx server hosts

The host environment is required to pass to its PMIx server library all non-environmental events that directly relate to a registered namespace without the PMIx server library explicitly requesting them. Environmental events are to be translated to their nearest PMIx equivalent code as defined in the range between **PMIX_EVENT_SYS_BASE** and **PMIX_EVENT_SYS_OTHER** (inclusive).

17.3.15 pmix_server_deregister_events_fn_t

Summary

Deregister to receive notifications for the specified events.

Format

C

```
typedef pmix_status_t (*pmix_server_deregister_events_fn_t) (  
    pmix_status_t *codes,  
    size_t ncodes,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN codes

Array of [pmix_status_t](#) values (array of handles)

IN ncodes

Number of elements in the *codes* array (integer)

IN cbfunc

Callback function [pmix_op_cbfunc_t](#) (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns one of the following:

- [PMIX_SUCCESS](#), indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function prior to returning from the API.
- [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- [PMIX_ERR_NOT_SUPPORTED](#), indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Description

Deregister to receive notifications for the specified events to which the PMIx server has previously registered.

Advice to PMIx library implementers

The PMIx server library must track all client registrations. This module function shall only be called when:

- the library is deregistering environmental codes (i.e., a PMIx codes in the range between [PMIX_EVENT_SYS_BASE](#) and [PMIX_EVENT_SYS_OTHER](#), inclusive) or codes that lies outside the defined PMIx range of constants; and

- no client (including the server library itself) remains registered for notifications on any included code - i.e., a code should be included in this call only when no registered notifications against it remain.

17.3.16 pmix_server_notify_event_fn_t

Summary

Notify the specified processes of an event.

Format

C

```
typedef pmix_status_t (*pmix_server_notify_event_fn_t) (  
    pmix_status_t code,  
    const pmix_proc_t *source,  
    pmix_data_range_t range,  
    pmix_info_t info[],  
    size_t ninfo,  
    pmix_op_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN code

The `pmix_status_t` event code being referenced structure (handle)

IN source

`pmix_proc_t` of process that generated the event (handle)

IN range

`pmix_data_range_t` range over which the event is to be distributed (handle)

IN info

Optional array of `pmix_info_t` structures containing additional information on the event (array of handles)

IN ninfo

Number of elements in the `info` array (integer)

IN cbfunc

Callback function `pmix_op_cbfunc_t` (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the host must not invoke the callback function prior to returning from the API.

- 1 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
2 returned *success* - the *cbfunc* will not be called
- 3 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
4 request, even though the function entry was provided in the server module - the *cbfunc* will not
5 be called
- 6 • a PMIx error constant indicating either an error in the input or that the request was immediately
7 processed and failed - the *cbfunc* will not be called

Required Attributes

8 PMIx libraries are required to pass any provided attributes to the host environment for processing.

9 Host environments that provide this module entry point are required to support the following
10 attributes:

11 **PMIX_RANGE** "`pmix.range`" (`pmix_data_range_t`)

12 Define constraints on the processes that can access published data or generated events or
13 define constraints on the provider of data when looking up published data.

14 **Description**

15 Notify the specified processes (described through a combination of *range* and attributes provided in
16 the *info* array) of an event generated either by the PMIx server itself or by one of its local clients.
17 The process generating the event is provided in the *source* parameter, and any further descriptive
18 information is included in the *info* array.

19 Note that the PMIx server library is not allowed to echo any event given to it by its host via the
20 **PMIx_Notify_event** API back to the host through the
21 **pmix_server_notify_event_fn_t** server module function.

Advice to PMIx server hosts

22 The callback function is to be executed once the host environment no longer requires that the PMIx
23 server library maintain the provided data structures. It does not necessarily indicate that the event
24 has been delivered to any process, nor that the event has been distributed for delivery

25 **17.3.17 pmix_server_listener_fn_t**

26 **Summary**

27 Register a socket the host server can monitor for connection requests.

Format

C

```
typedef pmix_status_t (*pmix_server_listener_fn_t) (  
    int listening_sd,  
    pmix_connection_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN `incoming_sd`

(integer)

IN `cbfunc`

Callback function `pmix_connection_cbfunc_t` (function reference)

IN `cbdata`

(memory reference)

Returns `PMIX_SUCCESS` indicating that the request is accepted, or a negative value corresponding to a PMIx error constant indicating that the request has been rejected.

Description

Register a socket the host environment can monitor for connection requests, harvest them, and then call the PMIx server library's internal callback function for further processing. A listener thread is essential to efficiently harvesting connection requests from large numbers of local clients such as occur when running on large SMPs. The host server listener is required to call `accept` on the incoming connection request, and then pass the resulting socket to the provided `cbfunc`. A `NULL` for this function will cause the internal PMIx server to spawn its own listener thread.

17.3.17.1 PMIx Client Connection Callback Function

Summary

Callback function for incoming connection request from a local client.

Format

C

```
typedef void (*pmix_connection_cbfunc_t) (  
    int incoming_sd, void *cbdata);
```

C

IN `incoming_sd`

(integer)

IN `cbdata`

(memory reference)

Description

Callback function for incoming connection requests from local clients - only used by host environments that wish to directly handle socket connection requests.

1 17.3.18 pmix_server_query_fn_t

2 Summary

3 Query information from the resource manager.

4 Format

PMIx v2.0

C

```
5 typedef pmix_status_t (*pmix_server_query_fn_t) (  
6     pmix_proc_t *proct,  
7     pmix_query_t *queries,  
8     size_t nqueries,  
9     pmix_info_cbfunc_t cbfunc,  
10    void *cbdata);
```

C

11 IN proct

12 [pmix_proc_t](#) structure of the requesting process (handle)

13 IN queries

14 Array of [pmix_query_t](#) structures (array of handles)

15 IN nqueries

16 Number of elements in the *queries* array (integer)

17 IN cbfunc

18 Callback function [pmix_info_cbfunc_t](#) (function reference)

19 IN cbdata

20 Data to be passed to the callback function (memory reference)

21 Returns one of the following:

- 22 • [PMIX_SUCCESS](#), indicating that the request is being processed by the host environment - result
23 will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function
24 prior to returning from the API.
- 25 • [PMIX_OPERATION_SUCCEEDED](#), indicating that the request was immediately processed and
26 returned *success* - the *cbfunc* will not be called
- 27 • [PMIX_ERR_NOT_SUPPORTED](#), indicating that the host environment does not support the
28 request, even though the function entry was provided in the server module - the *cbfunc* will not
29 be called
- 30 • a PMIx error constant indicating either an error in the input or that the request was immediately
31 processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (uint32_t)

Effective user ID of the connecting process.

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group ID of the connecting process.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_QUERY_NAMESPACES "pmix.qry.ns" (char*)

Request a comma-delimited list of active namespaces. NO QUALIFIERS.

PMIX_QUERY_JOB_STATUS "pmix.qry.jst" (pmix_status_t)

Status of a specified, currently executing job. REQUIRED QUALIFIER: **PMIX_NAMESPACE** indicating the namespace whose status is being queried.

PMIX_QUERY_QUEUE_LIST "pmix.qry.qlst" (char*)

Request a comma-delimited list of scheduler queues. NO QUALIFIERS.

PMIX_QUERY_QUEUE_STATUS "pmix.qry.qst" (char*)

Returns status of a specified scheduler queue, expressed as a string. OPTIONAL QUALIFIERS: **PMIX_ALLOC_QUEUE** naming specific queue whose status is being requested.

PMIX_QUERY_PROC_TABLE "pmix.qry.ptable" (char*)

Returns a (**pmix_data_array_t**) array of **pmix_proc_info_t**, one entry for each process in the specified namespace, ordered by process job rank. REQUIRED QUALIFIER: **PMIX_NAMESPACE** indicating the namespace whose process table is being queried.

PMIX_QUERY_LOCAL_PROC_TABLE "pmix.qry.lptable" (char*)

Returns a (**pmix_data_array_t**) array of **pmix_proc_info_t**, one entry for each process in the specified namespace executing on the same node as the requester, ordered by process job rank. REQUIRED QUALIFIER: **PMIX_NAMESPACE** indicating the namespace whose local process table is being queried. OPTIONAL QUALIFIER: **PMIX_HOSTNAME** indicating the host whose local process table is being queried. By default, the query assumes that the host upon which the request was made is to be used.

PMIX_QUERY_SPAWN_SUPPORT "pmix.qry.spawn" (bool)

Return a comma-delimited list of supported spawn attributes. NO QUALIFIERS.

PMIX_QUERY_DEBUG_SUPPORT "pmix.qry.debug" (bool)

Return a comma-delimited list of supported debug attributes. NO QUALIFIERS.

1 **PMIX_QUERY_MEMORY_USAGE** "pmix.qry.mem" (bool)
2 Return information on memory usage for the processes indicated in the qualifiers.
3 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** and **PMIX_RANK**, or **PMIX_PROCID** of
4 specific process(es) whose memory usage is being requested.

5 **PMIX_QUERY_LOCAL_ONLY** "pmix.qry.local" (bool)
6 Constrain the query to local information only. NO QUALIFIERS.

7 **PMIX_QUERY_REPORT_AVG** "pmix.qry.avg" (bool)
8 Report only average values for sampled information. NO QUALIFIERS.

9 **PMIX_QUERY_REPORT_MINMAX** "pmix.qry.minmax" (bool)
10 Report minimum and maximum values. NO QUALIFIERS.

11 **PMIX_QUERY_ALLOC_STATUS** "pmix.query.alloc" (char*)
12 String identifier of the allocation whose status is being requested. NO QUALIFIERS.

13 **PMIX_TIME_REMAINING** "pmix.time.remaining" (char*)
14 Query number of seconds (**uint32_t**) remaining in allocation for the specified namespace.
15 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** of the namespace whose info is being
16 requested (defaults to allocation containing the caller).

Description

Query information from the host environment. The query will include the namespace/rank of the process that is requesting the info, an array of **pmix_query_t** describing the request, and a callback function/data for the return.

Advice to PMIx library implementers

The PMIx server library should not block in this function as the host environment may, depending upon the information being requested, require significant time to respond.

17.3.19 pmix_server_tool_connection_fn_t

Summary

Register that a tool has connected to the server.

Format

C

```
typedef void (*pmix_server_tool_connection_fn_t) (  
    pmix_info_t info[], size_t ninfo,  
    pmix_tool_connection_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN `info`

Array of `pmix_info_t` structures (array of handles)

IN `ninfo`

Number of elements in the `info` array (integer)

IN `cbfunc`

Callback function `pmix_tool_connection_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Required Attributes

PMIx libraries are required to pass the following attributes in the `info` array:

PMIX_USERID "pmix.euid" (`uint32_t`)

Effective user ID of the connecting process.

PMIX_GRPID "pmix.egid" (`uint32_t`)

Effective group ID of the connecting process.

PMIX_TOOL_NAMESPACE "pmix.tool.namespace" (`char*`)

Name of the namespace to use for this tool. This must be included only if the tool already has an assigned namespace.

PMIX_TOOL_RANK "pmix.tool.rank" (`uint32_t`)

Rank of this tool. This must be included only if the tool already has an assigned rank.

PMIX_CREDENTIAL "pmix.cred" (`char*`)

Security credential assigned to the process.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_FWD_STDOUT "pmix.fwd.stdout" (bool)

Requests that the ability to forward the **stdout** of the spawned processes be maintained. The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and other options for delivery of the forwarded output.

PMIX_FWD_STDERR "pmix.fwd.stderr" (bool)

Requests that the ability to forward the **stderr** of the spawned processes be maintained. The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and other options for delivery of the forwarded output.

PMIX_FWD_STDIN "pmix.fwd.stdin" (pmix_rank_t)

The requester intends to push information from its **stdin** to the indicated process. The local spawn agent should, therefore, ensure that the **stdin** channel to that process remains available. A rank of **PMIX_RANK_WILDCARD** indicates that all processes in the spawned job are potential recipients. The requester will issue a call to **PMIx_IOF_push** to initiate the actual forwarding of information to specified targets - this attribute simply requests that the IL retain the ability to forward the information to the designated targets.

PMIX_VERSION_INFO "pmix.version" (char*)

PMIx version of the library being used by the connecting process.

Description

Register that a tool has connected to the server, possibly requesting that the tool be assigned a namespace/rank identifier for further interactions. The **pmix_info_t** array is used to pass qualifiers for the connection request, including the effective uid and gid of the calling tool for authentication purposes.

If the tool already has an assigned process identifier, then this must be indicated in the *info* array. The host is responsible for checking that the provided namespace does not conflict with any currently known assignments, returning an appropriate error in the callback function if a conflict is found.

The host environment is solely responsible for authenticating and authorizing the connection using whatever means it deems appropriate. If certificates or other authentication information are required, then the tool must provide them. The conclusion of those operations shall be communicated back to the PMIx server library via the callback function.

Approval or rejection of the connection request shall be returned in the *status* parameter of the **pmix_tool_connection_cbfunc_t**. If the connection is refused, the PMIx server library must terminate the connection attempt. The host must not execute the callback function prior to returning from the API.

1 17.3.19.1 Tool connection attributes

2 Attributes associated with tool connections.

3 **PMIX_USERID** "pmix.euid" (uint32_t)

4 Effective user ID of the connecting process.

5 **PMIX_GRPID** "pmix.egid" (uint32_t)

6 Effective group ID of the connecting process.

7 **PMIX_VERSION_INFO** "pmix.version" (char*)

8 PMIx version of the library being used by the connecting process.

9 17.3.19.2 PMIx Tool Connection Callback Function

10 Summary

11 Callback function for incoming tool connections.

12 Format

PMIx v2.0

```
typedef void (*pmix_tool_connection_cbfunc_t) (  
    pmix_status_t status,  
    pmix_proc_t *proc, void *cbdata);
```

16 **IN status**

[pmix_status_t](#) value (handle)

18 **IN proc**

[pmix_proc_t](#) structure containing the identifier assigned to the tool (handle)

20 **IN cbdata**

Data to be passed (memory reference)

22 Description

23 Callback function for incoming tool connections. The host environment shall provide a
24 namespace/rank identifier for the connecting tool.

Advice to PMIx server hosts

25 It is assumed that **rank=0** will be the normal assignment, but allow for the future possibility of a
26 parallel set of tools connecting, and thus each process requiring a unique rank.

27 17.3.20 pmix_server_log_fn_t

28 Summary

29 Log data on behalf of a client.

Format

C

```
typedef void (*pmix_server_log_fn_t) (  
    const pmix_proc_t *client,  
    const pmix_info_t data[], size_t ndata,  
    const pmix_info_t directives[], size_t ndirs,  
    pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

IN client
pmix_proc_t structure (handle)

IN data
Array of info structures (array of handles)

IN ndata
Number of elements in the *data* array (integer)

IN directives
Array of info structures (array of handles)

IN ndirs
Number of elements in the *directives* array (integer)

IN cbfunc
Callback function pmix_op_cbfunc_t (function reference)

IN cbdata
Data to be passed to the callback function (memory reference)

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (uint32_t)
Effective user ID of the connecting process.

PMIX_GRPID "pmix.egid" (uint32_t)
Effective group ID of the connecting process.

Host environments that provide this module entry point are required to support the following attributes:

PMIX_LOG_STDERR "pmix.log.stderr" (char*)
Log string to *stderr*.

PMIX_LOG_STDOUT "pmix.log.stdout" (char*)
Log string to *stdout*.

PMIX_LOG_SYSLOG "pmix.log.syslog" (char*)

1 Log data to syslog. Defaults to **ERROR** priority. Will log to global syslog if available,
2 otherwise to local syslog.

Optional Attributes

3 The following attributes are optional for host environments that support this operation:

4 **PMIX_LOG_MSG** "pmix.log.msg" (pmix_byte_object_t)

5 Message blob to be sent somewhere.

6 **PMIX_LOG_EMAIL** "pmix.log.email" (pmix_data_array_t)

7 Log via email based on **pmix_info_t** containing directives.

8 **PMIX_LOG_EMAIL_ADDR** "pmix.log.emaddr" (char*)

9 Comma-delimited list of email addresses that are to receive the message.

10 **PMIX_LOG_EMAIL_SUBJECT** "pmix.log.emsub" (char*)

11 Subject line for email.

12 **PMIX_LOG_EMAIL_MSG** "pmix.log.emmsg" (char*)

13 Message to be included in email.

Description

14 Log data on behalf of a client. This function is not intended for output of computational results, but
15 rather for reporting status and error messages. The host must not execute the callback function prior
16 to returning from the API.
17

18 17.3.21 pmix_server_alloc_fn_t

19 Summary

20 Request allocation operations on behalf of a client.

Format

```
typedef pmix_status_t (*pmix_server_alloc_fn_t) (  
    const pmix_proc_t *client,  
    pmix_alloc_directive_t directive,  
    const pmix_info_t data[],  
    size_t ndata,  
    pmix_info_cbfunc_t cbfunc,  
    void *cbdata);
```

IN client

`pmix_proc_t` structure of process making request (handle)

IN directive

Specific action being requested (`pmix_alloc_directive_t`)

IN data

Array of info structures (array of handles)

IN ndata

Number of elements in the *data* array (integer)

IN cbfunc

Callback function `pmix_info_cbfunc_t` (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

PMIx libraries are required to pass any provided attributes to the host environment for processing. In addition, the following attributes are required to be included in the passed *info* array:

PMIX_USERID "pmix.euid" (`uint32_t`)

Effective user ID of the connecting process.

1 **PMIX_GRPID** "pmix.egid" (uint32_t)
2 Effective group ID of the connecting process.

3
4 Host environments that provide this module entry point are required to support the following
5 attributes:

6 **PMIX_ALLOC_ID** "pmix.alloc.id" (char*)
7 A string identifier (provided by the host environment) for the resulting allocation which can
8 later be used to reference the allocated resources in, for example, a call to **PMIx_Spawn**.

9 **PMIX_ALLOC_NUM_NODES** "pmix.alloc.nnodes" (uint64_t)
10 The number of nodes being requested in an allocation request.

11 **PMIX_ALLOC_NUM_CPUS** "pmix.alloc.ncpus" (uint64_t)
12 Number of PUs being requested in an allocation request.

13 **PMIX_ALLOC_TIME** "pmix.alloc.time" (uint32_t)
14 Total session time (in seconds) being requested in an allocation request.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

15 The following attributes are optional for host environments that support this operation:

16 **PMIX_ALLOC_NODE_LIST** "pmix.alloc.nlist" (char*)
17 Regular expression of the specific nodes being requested in an allocation request.

18 **PMIX_ALLOC_NUM_CPU_LIST** "pmix.alloc.ncpulist" (char*)
19 Regular expression of the number of PUs for each node being requested in an allocation
20 request.

21 **PMIX_ALLOC_CPU_LIST** "pmix.alloc.cpulist" (char*)
22 Regular expression of the specific PUs being requested in an allocation request.

23 **PMIX_ALLOC_MEM_SIZE** "pmix.alloc.msize" (float)
24 Number of Megabytes[base2] of memory (per process) being requested in an allocation
25 request.

26 **PMIX_ALLOC_FABRIC** "pmix.alloc.net" (array)
27 Array of **pmix_info_t** describing requested fabric resources. This must include at least:
28 **PMIX_ALLOC_FABRIC_ID**, **PMIX_ALLOC_FABRIC_TYPE**, and
29 **PMIX_ALLOC_FABRIC_ENDPTS**, plus whatever other descriptors are desired.

30 **PMIX_ALLOC_FABRIC_ID** "pmix.alloc.netid" (char*)

1 The key to be used when accessing this requested fabric allocation. The fabric allocation
2 will be returned/stored as a `pmix_data_array_t` of `pmix_info_t` whose first
3 element is composed of this key and the allocated resource description. The type of the
4 included value depends upon the fabric support. For example, a TCP allocation might
5 consist of a comma-delimited string of socket ranges such as "32000–32100,
6 33005, 38123–38146". Additional array entries will consist of any provided resource
7 request directives, along with their assigned values. Examples include:
8 `PMIX_ALLOC_FABRIC_TYPE` - the type of resources provided;
9 `PMIX_ALLOC_FABRIC_PLANE` - if applicable, what plane the resources were assigned
10 from; `PMIX_ALLOC_FABRIC_QOS` - the assigned QoS; `PMIX_ALLOC_BANDWIDTH` -
11 the allocated bandwidth; `PMIX_ALLOC_FABRIC_SEC_KEY` - a security key for the
12 requested fabric allocation. NOTE: the array contents may differ from those requested,
13 especially if `PMIX_INFO_REQD` was not set in the request.

14 `PMIX_ALLOC_BANDWIDTH` "pmix.alloc.bw" (float)

15 Fabric bandwidth (in Megabits[base2]/sec) for the job being requested in an allocation
16 request.

17 `PMIX_ALLOC_FABRIC_QOS` "pmix.alloc.netqos" (char*)

18 Fabric quality of service level for the job being requested in an allocation request.



19 Description

20 Request new allocation or modifications to an existing allocation on behalf of a client. Several
21 broad categories are envisioned, including the ability to:

- 22 • Request allocation of additional resources, including memory, bandwidth, and compute for an
23 existing allocation. Any additional allocated resources will be considered as part of the current
24 allocation, and thus will be released at the same time.
- 25 • Request a new allocation of resources. Note that the new allocation will be disjoint from (i.e., not
26 affiliated with) the allocation of the requestor - thus the termination of one allocation will not
27 impact the other.
- 28 • Extend the reservation on currently allocated resources, subject to scheduling availability and
29 priorities.
- 30 • Return no-longer-required resources to the scheduler. This includes the *loan* of resources back to
31 the scheduler with a promise to return them upon subsequent request.

32 The callback function provides a *status* to indicate whether or not the request was granted, and to
33 provide some information as to the reason for any denial in the `pmix_info_cbfunc_t` array of
34 `pmix_info_t` structures.

35 17.3.22 pmix_server_job_control_fn_t

36 Summary

37 Execute a job control action on behalf of a client.

Format

C

```
typedef pmix_status_t (*pmix_server_job_control_fn_t) (  
    const pmix_proc_t *requestor,  
    const pmix_proc_t targets[],  
    size_t ntargets,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    pmix_info_cbfunc_t cbfunc,  
    void *cbdata);
```

C

- IN requestor**
[pmix_proc_t](#) structure of requesting process (handle)
- IN targets**
Array of proc structures (array of handles)
- IN ntargets**
Number of elements in the *targets* array (integer)
- IN directives**
Array of info structures (array of handles)
- IN ndirs**
Number of elements in the *info* array (integer)
- IN cbfunc**
Callback function [pmix_info_cbfunc_t](#) (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

1 PMIx libraries are required to pass any attributes provided by the client to the host environment for
2 processing. In addition, the following attributes are required to be included in the passed *info* array:

3 **PMIX_USERID** "pmix.euid" (uint32_t)

4 Effective user ID of the connecting process.

5 **PMIX_GRPID** "pmix.egid" (uint32_t)

6 Effective group ID of the connecting process.

8 Host environments that provide this module entry point are required to support the following
9 attributes:

10 **PMIX_JOB_CTRL_ID** "pmix.jctrl.id" (char*)

11 Provide a string identifier for this request. The user can provide an identifier for the
12 requested operation, thus allowing them to later request status of the operation or to
13 terminate it. The host, therefore, shall track it with the request for future reference.

14 **PMIX_JOB_CTRL_PAUSE** "pmix.jctrl.pause" (bool)

15 Pause the specified processes.

16 **PMIX_JOB_CTRL_RESUME** "pmix.jctrl.resume" (bool)

17 Resume ("un-pause") the specified processes.

18 **PMIX_JOB_CTRL_KILL** "pmix.jctrl.kill" (bool)

19 Forcibly terminate the specified processes and cleanup.

20 **PMIX_JOB_CTRL_SIGNAL** "pmix.jctrl.sig" (int)

21 Send given signal to specified processes.

22 **PMIX_JOB_CTRL_TERMINATE** "pmix.jctrl.term" (bool)

23 Politely terminate the specified processes.

Optional Attributes

24 The following attributes are optional for host environments that support this operation:

25 **PMIX_JOB_CTRL_CANCEL** "pmix.jctrl.cancel" (char*)

26 Cancel the specified request - the provided request ID must match the
27 **PMIX_JOB_CTRL_ID** provided to a previous call to **PMIx_Job_control**. An ID of
28 **NULL** implies cancel all requests from this requestor.

29 **PMIX_JOB_CTRL_RESTART** "pmix.jctrl.restart" (char*)

30 Restart the specified processes using the given checkpoint ID.

31 **PMIX_JOB_CTRL_CHECKPOINT** "pmix.jctrl.ckpt" (char*)

32 Checkpoint the specified processes and assign the given ID to it.

1 **PMIX_JOB_CTRL_CHECKPOINT_EVENT** "pmix.jctrl.ckptev" (bool)
2 Use event notification to trigger a process checkpoint.

3 **PMIX_JOB_CTRL_CHECKPOINT_SIGNAL** "pmix.jctrl.ckptsig" (int)
4 Use the given signal to trigger a process checkpoint.

5 **PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT** "pmix.jctrl.ckptsig" (int)
6 Time in seconds to wait for a checkpoint to complete.

7 **PMIX_JOB_CTRL_CHECKPOINT_METHOD**
8 "pmix.jctrl.ckmethod" (pmix_data_array_t)
9 Array of **pmix_info_t** declaring each method and value supported by this application.

10 **PMIX_JOB_CTRL_PROVISION** "pmix.jctrl.pvn" (char*)
11 Regular expression identifying nodes that are to be provisioned.

12 **PMIX_JOB_CTRL_PROVISION_IMAGE** "pmix.jctrl.pvning" (char*)
13 Name of the image that is to be provisioned.

14 **PMIX_JOB_CTRL_PREEMPTIBLE** "pmix.jctrl.preempt" (bool)
15 Indicate that the job can be pre-empted.

▲-----▲

16 Description

17 Execute a job control action on behalf of a client. The *targets* array identifies the processes to
18 which the requested job control action is to be applied. A **NULL** value can be used to indicate all
19 processes in the caller's namespace. The use of **PMIX_RANK_WILDCARD** can also be used to
20 indicate that all processes in the given namespace are to be included.

21 The directives are provided as **pmix_info_t** structures in the *directives* array. The callback
22 function provides a *status* to indicate whether or not the request was granted, and to provide some
23 information as to the reason for any denial in the **pmix_info_cbfnc_t** array of
24 **pmix_info_t** structures.

25 17.3.23 pmix_server_monitor_fn_t

26 Summary

27 Request that a client be monitored for activity.

Format

C

```
typedef pmix_status_t (*pmix_server_monitor_fn_t) (  
    const pmix_proc_t *requestor,  
    const pmix_info_t *monitor,  
    pmix_status_t error,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    pmix_info_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN requestor

[pmix_proc_t](#) structure of requesting process (handle)

IN monitor

[pmix_info_t](#) identifying the type of monitor being requested (handle)

IN error

Status code to use in generating event if alarm triggers (integer)

IN directives

Array of info structures (array of handles)

IN ndirs

Number of elements in the *info* array (integer)

IN cbfunc

Callback function [pmix_info_cbfunc_t](#) (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the host must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

This entry point is only called for monitoring requests that are not directly supported by the PMIx server library itself.

Required Attributes

1 If supported by the PMix server library, then the library must not pass any supported attributes to
2 the host environment. Any attributes provided by the client that are not directly supported by the
3 server library must be passed to the host environment if it provides this module entry. In addition,
4 the following attributes are required to be included in the passed *info* array:

5 **PMIX_USERID** "pmix.euid" (uint32_t)

6 Effective user ID of the connecting process.

7 **PMIX_GRPID** "pmix.egid" (uint32_t)

8 Effective group ID of the connecting process.

9 Host environments are not required to support any specific monitoring attributes.

Optional Attributes

10 The following attributes may be implemented by a host environment.

11 **PMIX_MONITOR_ID** "pmix.monitor.id" (char*)

12 Provide a string identifier for this request.

13 **PMIX_MONITOR_CANCEL** "pmix.monitor.cancel" (char*)

14 Identifier to be canceled (NULL means cancel all monitoring for this process).

15 **PMIX_MONITOR_APP_CONTROL** "pmix.monitor.appctrl" (bool)

16 The application desires to control the response to a monitoring event - i.e., the application is
17 requesting that the host environment not take immediate action in response to the event (e.g.,
18 terminating the job).

19 **PMIX_MONITOR_HEARTBEAT** "pmix.monitor.mbeat" (void)

20 Register to have the PMix server monitor the requestor for heartbeats.

21 **PMIX_MONITOR_HEARTBEAT_TIME** "pmix.monitor.btime" (uint32_t)

22 Time in seconds before declaring heartbeat missed.

23 **PMIX_MONITOR_HEARTBEAT_DROPS** "pmix.monitor.bdrop" (uint32_t)

24 Number of heartbeats that can be missed before generating the event.

25 **PMIX_MONITOR_FILE** "pmix.monitor.fmon" (char*)

26 Register to monitor file for signs of life.

27 **PMIX_MONITOR_FILE_SIZE** "pmix.monitor.fsize" (bool)

28 Monitor size of given file is growing to determine if the application is running.

29 **PMIX_MONITOR_FILE_ACCESS** "pmix.monitor.faccess" (char*)

30 Monitor time since last access of given file to determine if the application is running.

31 **PMIX_MONITOR_FILE_MODIFY** "pmix.monitor.fmod" (char*)

32 Monitor time since last modified of given file to determine if the application is running.

1 **PMIX_MONITOR_FILE_CHECK_TIME** "pmix.monitor.ftime" (uint32_t)

2 Time in seconds between checking the file.

3 **PMIX_MONITOR_FILE_DROPS** "pmix.monitor.fdrop" (uint32_t)

4 Number of file checks that can be missed before generating the event.



5 Description

6 Request that a client be monitored for activity.

7 17.3.24 pmix_server_get_cred_fn_t

8 Summary

9 Request a credential from the host environment.

10 Format

PMIx v3.0

C

```
11 typedef pmix_status_t (*pmix_server_get_cred_fn_t) (  
12     const pmix_proc_t *proc,  
13     const pmix_info_t directives[],  
14     size_t ndirs,  
15     pmix_credential_cbfunc_t cbfunc,  
16     void *cbdata);
```

C

17 **IN proc**

18 **pmix_proc_t** structure of requesting process (handle)

19 **IN directives**

20 Array of info structures (array of handles)

21 **IN ndirs**

22 Number of elements in the *info* array (integer)

23 **IN cbfunc**

24 Callback function to return the credential (**pmix_credential_cbfunc_t** function
25 reference)

26 **IN cbdata**

27 Data to be passed to the callback function (memory reference)

- 28 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
29 will be returned in the provided *cbfunc*
- 30 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
31 request, even though the function entry was provided in the server module - the *cbfunc* will not
32 be called
- 33 • a PMIx error constant indicating either an error in the input or that the request was immediately
34 processed and failed - the *cbfunc* will not be called

Required Attributes

If the PMIx library does not itself provide the requested credential, then it is required to pass any attributes provided by the client to the host environment for processing. In addition, it must include the following attributes in the passed *info* array:

PMIX_USERID "pmix.euid" (uint32_t)

Effective user ID of the connecting process.

PMIX_GRPID "pmix.egid" (uint32_t)

Effective group ID of the connecting process.

Optional Attributes

The following attributes are optional for host environments that support this operation:

PMIX_CRED_TYPE "pmix.sec.ctype" (char*)

When passed in **PMIx_Get_credential**, a prioritized, comma-delimited list of desired credential types for use in environments where multiple authentication mechanisms may be available. When returned in a callback function, a string identifier of the credential type.

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Request a credential from the host environment.

17.3.24.1 Credential callback function

Summary

Callback function to return a requested security credential

Format

C

```
typedef void (*pmix_credential_cfunc_t) (  
    pmix_status_t status,  
    pmix_byte_object_t *credential,  
    pmix_info_t info[], size_t ninfo,  
    void *cbdata);
```

C

IN status

[pmix_status_t](#) value (handle)

IN credential

[pmix_byte_object_t](#) structure containing the security credential (handle)

IN info

Array of provided by the system to pass any additional information about the credential - e.g., the identity of the issuing agent. (handle)

IN ninfo

Number of elements in *info* ([size_t](#))

IN cbdata

Object passed in original request (memory reference)

Description

Define a callback function to return a requested security credential. Information provided by the issuing agent can subsequently be used by the application for a variety of purposes. Examples include:

- checking identified authorizations to determine what requests/operations are feasible as a means to steering [workflows](#)
- compare the credential type to that of the local SMS for compatibility

Advice to users

The credential is opaque and therefore understandable only by a service compatible with the issuer. The *info* array is owned by the PMIx library and is not to be released or altered by the receiving party.

17.3.25 pmix_server_validate_cred_fn_t

Summary

Request validation of a credential.

Format

```
typedef pmix_status_t (*pmix_server_validate_cred_fn_t) (  
    const pmix_proc_t *proc,  
    const pmix_byte_object_t *cred,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    pmix_validation_cbfunc_t cbfunc,  
    void *cbdata);
```

IN `proc`

`pmix_proc_t` structure of requesting process (handle)

IN `cred`

Pointer to `pmix_byte_object_t` containing the credential (handle)

IN `directives`

Array of info structures (array of handles)

IN `ndirs`

Number of elements in the *info* array (integer)

IN `cbfunc`

Callback function to return the result (`pmix_validation_cbfunc_t` function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

If the PMIx library does not itself validate the credential, then it is required to pass any attributes provided by the client to the host environment for processing. In addition, it must include the following attributes in the passed *info* array:

PMIX_USERID "pmix.euid" (`uint32_t`)

Effective user ID of the connecting process.

1 **PMIX_GRPID** "pmix.egid" (uint32_t)
2 Effective group ID of the connecting process.

3
4 Host environments are not required to support any specific attributes.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

5 The following attributes are optional for host environments that support this operation:

6 **PMIX_TIMEOUT** "pmix.timeout" (int)
7 Time in seconds before the specified operation should time out (zero indicating infinite) and
8 return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions
9 caused by multiple layers (client, server, and host) simultaneously timing the operation.

▲-----▲

10 **Description**

11 Request validation of a credential obtained from the host environment via a prior call to the
12 **pmix_server_get_cred_fn_t** module entry.

13 **17.3.26 Credential validation callback function**

14 **Summary**

15 Callback function for security credential validation.

Format

C

```
typedef void (*pmix_validation_cbfunc_t) (  
    pmix_status_t status,  
    pmix_info_t info[], size_t ninfo,  
    void *cbdata);
```

C

IN status

[pmix_status_t](#) value (handle)

IN info

Array of [pmix_info_t](#) provided by the system to pass any additional information about the authentication - e.g., the effective userid and group id of the certificate holder, and any related authorizations (handle)

IN ninfo

Number of elements in *info* ([size_t](#))

IN cbdata

Object passed in original request (memory reference)

The returned status shall be one of the following:

- [PMIX_SUCCESS](#), indicating that the request was processed and returned *success* (i.e., the credential was both valid and any information it contained was successfully processed). Details of the result will be returned in the *info* array
- a PMIx error constant indicating either an error in the parsing of the credential or that the request was refused

Description

Define a validation callback function to indicate if a provided credential is valid, and any corresponding information regarding authorizations and other security matters.

Advice to users

The precise contents of the array will depend on the host environment and its associated security system. At the minimum, it is expected (but not required) that the array will contain entries for the [PMIX_USERID](#) and [PMIX_GRPID](#) of the client described in the credential. The *info* array is owned by the PMIx library and is not to be released or altered by the receiving party.

17.3.27 pmix_server_iof_fn_t

Summary

Request the specified IO channels be forwarded from the given array of processes.

Format

```
typedef pmix_status_t (*pmix_server_iof_fn_t) (  
    const pmix_proc_t procs[],  
    size_t nprocs,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    pmix_iof_channel_t channels,  
    pmix_op_cbfunc_t cbfunc, void *cbdata);
```

IN `procs`

Array `pmix_proc_t` identifiers whose IO is being requested (handle)

IN `nprocs`

Number of elements in `procs` (`size_t`)

IN `directives`

Array of `pmix_info_t` structures further defining the request (array of handles)

IN `ndirs`

Number of elements in the `info` array (integer)

IN `channels`

Bitmask identifying the channels to be forwarded (`pmix_iof_channel_t`)

IN `cbfunc`

Callback function `pmix_op_cbfunc_t` (function reference)

IN `cbdata`

Data to be passed to the callback function (memory reference)

Returns one of the following:

- `PMIX_SUCCESS`, indicating that the request is being processed by the host environment - result will be returned in the provided `cbfunc`. Note that the library must not invoke the callback function prior to returning from the API.
- `PMIX_OPERATION_SUCCEEDED`, indicating that the request was immediately processed and returned *success* - the `cbfunc` will not be called
- `PMIX_ERR_NOT_SUPPORTED`, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the `cbfunc` will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the `cbfunc` will not be called

Required Attributes

The following attributes are required to be included in the passed `info` array:

`PMIX_USERID` "pmix.euid" (`uint32_t`)

1 Effective user ID of the connecting process.

2 **PMIX_GRPID** "pmix.egid" (uint32_t)

3 Effective group ID of the connecting process.

4

5 Host environments that provide this module entry point are required to support the following
6 attributes:

7 **PMIX_IOF_CACHE_SIZE** "pmix.iof.csize" (uint32_t)

8 The requested size of the PMIx server cache in bytes for each specified channel. By default,
9 the server is allowed (but not required) to drop all bytes received beyond the max size.

10 **PMIX_IOF_DROP_OLDEST** "pmix.iof.old" (bool)

11 In an overflow situation, the PMIx server is to drop the oldest bytes to make room in the
12 cache.

13 **PMIX_IOF_DROP_NEWEST** "pmix.iof.new" (bool)

14 In an overflow situation, the PMIx server is to drop any new bytes received until room
15 becomes available in the cache (default).



▼----- Optional Attributes -----▼

16 The following attributes may be supported by a host environment.

17 **PMIX_IOF_BUFFERING_SIZE** "pmix.iof.bsize" (uint32_t)

18 Requests that IO on the specified channel(s) be aggregated in the PMIx tool library until the
19 specified number of bytes is collected to avoid being called every time a block of IO arrives.
20 The PMIx tool library will execute the callback and reset the collection counter whenever the
21 specified number of bytes becomes available. Any remaining buffered data will be *flushed* to
22 the callback upon a call to deregister the respective channel.

23 **PMIX_IOF_BUFFERING_TIME** "pmix.iof.btime" (uint32_t)

24 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
25 size, this prevents IO from being held indefinitely while waiting for another payload to
26 arrive.



Description

Request the specified IO channels be forwarded from the given array of processes. An error shall be returned in the callback function if the requested service from any of the requested processes cannot be provided.

Advice to PMIx library implementers

The forwarding of stdin is a *push* process - processes cannot request that it be *pulled* from some other source. Requests including the `PMIX_FWD_STDIN_CHANNEL` channel will return a `PMIX_ERR_NOT_SUPPORTED` error.

17.3.27.1 IOF delivery function

Summary

Callback function for delivering forwarded IO to a process.

Format

PMIx v3.0

C

```
typedef void (*pmix_iof_cfunc_t) (  
    size_t iofhdlr, pmix_iof_channel_t channel,  
    pmix_proc_t *source, char *payload,  
    pmix_info_t info[], size_t ninfo);
```

C

IN iofhdlr

Registration number of the handler being invoked (`size_t`)

IN channel

bitmask identifying the channel the data arrived on (`pmix_iof_channel_t`)

IN source

Pointer to a `pmix_proc_t` identifying the namespace/rank of the process that generated the data (`char*`)

IN payload

Pointer to character array containing the data.

IN info

Array of `pmix_info_t` provided by the source containing metadata about the payload. This could include `PMIX_IOF_COMPLETE` (handle)

IN ninfo

Number of elements in *info* (`size_t`)

Description

Define a callback function for delivering forwarded IO to a process. This function will be called whenever data becomes available, or a specified buffering size and/or time has been met.

Advice to users

Multiple strings may be included in a given *payload*, and the *payload* may *not* be **NULL** terminated. The user is responsible for releasing the *payload* memory. The *info* array is owned by the PMIx library and is not to be released or altered by the receiving party.

17.3.28 pmix_server_stdin_fn_t

Summary

Pass standard input data to the host environment for transmission to specified recipients.

Format

PMIx v3.0

C

```
typedef pmix_status_t (*pmix_server_stdin_fn_t) (  
    const pmix_proc_t *source,  
    const pmix_proc_t targets[],  
    size_t ntargets,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    const pmix_byte_object_t *bo,  
    pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

- IN source**
[pmix_proc_t](#) structure of source process (handle)
- IN targets**
Array of [pmix_proc_t](#) target identifiers (handle)
- IN ntargets**
Number of elements in the *targets* array (integer)
- IN directives**
Array of info structures (array of handles)
- IN ndirs**
Number of elements in the *info* array (integer)
- IN bo**
Pointer to [pmix_byte_object_t](#) containing the payload (handle)
- IN cbfunc**
Callback function [pmix_op_cbfunc_t](#) (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

1 Returns one of the following:

- 2 • **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result
3 will be returned in the provided *cbfunc*. Note that the library must not invoke the callback
4 function prior to returning from the API.
- 5 • **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and
6 returned *success* - the *cbfunc* will not be called
- 7 • **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the
8 request, even though the function entry was provided in the server module - the *cbfunc* will not
9 be called
- 10 • a PMIx error constant indicating either an error in the input or that the request was immediately
11 processed and failed - the *cbfunc* will not be called

Required Attributes

12 The following attributes are required to be included in the passed *info* array:

13 **PMIX_USERID** "`pmix.euid`" (`uint32_t`)

14 Effective user ID of the connecting process.

15 **PMIX_GRPID** "`pmix.egid`" (`uint32_t`)

16 Effective group ID of the connecting process.

Description

17 Passes stdin to the host environment for transmission to specified recipients. The host environment
18 is responsible for forwarding the data to all locations that host the specified *targets* and delivering
19 the payload to the PMIx server library connected to those clients.
20

21 17.3.29 `pmix_server_grp_fn_t`

22 Summary

23 Request group operations (construct, destruct, etc.) on behalf of a set of processes.

Format

C

```
typedef pmix_status_t (*pmix_server_grp_fn_t) (  
    pmix_group_operation_t op,  
    char grp[],  
    const pmix_proc_t procs[],  
    size_t nprocs,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    pmix_info_cbfunc_t cbfunc,  
    void *cbdata);
```

C

- IN op**
[pmix_group_operation_t](#) value indicating operation the host is requested to perform (integer)
- IN grp**
Character string identifying the group (string)
- IN procs**
Array of [pmix_proc_t](#) identifiers of participants (handle)
- IN nprocs**
Number of elements in the *procs* array (integer)
- IN directives**
Array of info structures (array of handles)
- IN ndirs**
Number of elements in the *info* array (integer)
- IN cbfunc**
Callback function [pmix_info_cbfunc_t](#) (function reference)
- IN cbdata**
Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

The following attributes are required to be supported by a host environment.

PMIX_LOCAL_COLLECTIVE_STATUS "pmix.loc.col.st" (pmix_status_t)

Status code for local collective operation being reported to the host by the server library.

PMIX servers may aggregate the participation by local client processes in a collective operation - e.g., instead of passing individual client calls to **PMIX_Fence** up to the host environment, the server may pass only a single call to the host when all local participants have executed their **PMIX_Fence** call, thereby reducing the burden placed on the host.

However, in cases where the operation locally fails (e.g., if a participating client abnormally terminates prior to calling the operation), the server upcall functions to the host do not include a **pmix_status_t** by which the PMIX server can alert the host to that failure.

This attribute resolves that problem by allowing the server to pass the status information regarding the local collective operation.

Optional Attributes

The following attributes may be supported by a host environment.

PMIX_GROUP_ASSIGN_CONTEXT_ID "pmix.grp.actxid" (bool)

Requests that the RM assign a new context identifier to the newly created group. The identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range specified in the request. Thus, the value serves as a means of identifying the group within that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

PMIX_GROUP_LOCAL_ONLY "pmix.grp.lcl" (bool)

Group operation only involves local processes. PMIX implementations are *required* to automatically scan an array of group members for local vs remote processes - if only local processes are detected, the implementation need not execute a global collective for the operation unless a context ID has been requested from the host environment. This can result in significant time savings. This attribute can be used to optimize the operation by indicating whether or not only local processes are represented, thus allowing the implementation to bypass the scan.

PMIX_GROUP_ENDPT_DATA "pmix.grp.endpt" (pmix_byte_object_t)

Data collected during group construction to ensure communication between group members is supported upon completion of the operation.

PMIX_GROUP_OPTIONAL "pmix.grp.opt" (bool)

Participation is optional - do not return an error if any of the specified processes terminate without having joined. The default is **false**.

PMIX_RANGE "pmix.range" (pmix_data_range_t)

Define constraints on the processes that can access published data or generated events or define constraints on the provider of data when looking up published data.

1 The following attributes may be included in the host's response:

2 **PMIX_GROUP_ID** "pmix.grp.id" (char*)

3 User-provided group identifier - as the group identifier may be used in PMIx operations, the
4 user is required to ensure that the provided ID is unique within the scope of the host
5 environment (e.g., by including some user-specific or application-specific prefix or suffix to
6 the string).

7 **PMIX_GROUP_MEMBERSHIP** "pmix.grp.mbrs" (pmix_data_array_t*)

8 Array **pmix_proc_t** identifiers identifying the members of the specified group.

9 **PMIX_GROUP_CONTEXT_ID** "pmix.grp.ctxid" (size_t)

10 Context identifier assigned to the group by the host RM.

11 **PMIX_GROUP_ENDPT_DATA** "pmix.grp.endpt" (pmix_byte_object_t)

12 Data collected during group construction to ensure communication between group members
13 is supported upon completion of the operation.



14 Description

15 Perform the specified operation across the identified processes, plus any special actions included in
16 the directives. Return the result of any special action requests in the callback function when the
17 operation is completed. Actions may include a request (**PMIX_GROUP_ASSIGN_CONTEXT_ID**)
18 that the host assign a unique numerical (size_t) ID to this group - if given, the **PMIX_RANGE**
19 attribute will specify the range across which the ID must be unique (default to
20 **PMIX_RANGE_SESSION**).

21 17.3.29.1 Group Operation Constants

22 *PMIx v4.0*

The **pmix_group_operation_t** structure is a **uint8_t** value for specifying group
23 operations. All values were originally defined in version 4 of the standard unless otherwise marked.

24 **PMIX_GROUP_CONSTRUCT** 0 Construct a group composed of the specified processes - used
25 by a PMIx server library to direct host operation.

26 **PMIX_GROUP_DESTRUCT** 1 Destruct the specified group - used by a PMIx server library to
27 direct host operation.

28 17.3.30 pmix_server_fabric_fn_t

29 Summary

30 Request fabric-related operations (e.g., information on a fabric) on behalf of a tool or other process.

Format

C

```
typedef pmix_status_t (*pmix_server_fabric_fn_t) (  
    const pmix_proc_t *requestor,  
    pmix_fabric_operation_t op,  
    const pmix_info_t directives[],  
    size_t ndirs,  
    pmix_info_cbfunc_t cbfunc,  
    void *cbdata);
```

C

IN requestor

`pmix_proc_t` identifying the requestor (handle)

IN op

`pmix_fabric_operation_t` value indicating operation the host is requested to perform (integer)

IN directives

Array of info structures (array of handles)

IN ndirs

Number of elements in the *info* array (integer)

IN cbfunc

Callback function `pmix_info_cbfunc_t` (function reference)

IN cbdata

Data to be passed to the callback function (memory reference)

Returns one of the following:

- **PMIX_SUCCESS**, indicating that the request is being processed by the host environment - result will be returned in the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning from the API.
- **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed and returned *success* - the *cbfunc* will not be called
- **PMIX_ERR_NOT_SUPPORTED**, indicating that the host environment does not support the request, even though the function entry was provided in the server module - the *cbfunc* will not be called
- a PMIx error constant indicating either an error in the input or that the request was immediately processed and failed - the *cbfunc* will not be called

Required Attributes

The following directives are required to be supported by all hosts to aid users in identifying the fabric and (if applicable) the device to whom the operation references:

PMIX_FABRIC_VENDOR "pmix.fab.vndr" (**string**)

Name of the vendor (e.g., Amazon, Mellanox, HPE, Intel) for the specified fabric.

1 **PMIX_FABRIC_IDENTIFIER** "pmix.fab.id" (string)

2 An identifier for the specified fabric (e.g., MgmtEthernet, Slingshot-11, OmniPath-1).

3 **PMIX_FABRIC_PLANE** "pmix.fab.plane" (string)

4 ID string of a fabric plane (e.g., CIDR for Ethernet). When used as a modifier in a request
5 for information, specifies the plane whose information is to be returned. When used directly
6 as a key in a request, returns a **pmix_data_array_t** of string identifiers for all fabric
7 planes in the overall system.

8 **PMIX_FABRIC_DEVICE_INDEX** "pmix.fabdev.idx" (uint32_t)

9 Index of the device within an associated communication cost matrix.



10 Description

11 Perform the specified operation. Return the result of any requests in the callback function when the
12 operation is completed. Operations may, for example, include a request for fabric information. See
13 **pmix_fabric_t** for a list of expected information to be included in the response. Note that
14 requests for device index are to be returned in the callback function's array of **pmix_info_t**
15 using the **PMIX_FABRIC_DEVICE_INDEX** attribute.

CHAPTER 18

Tools and Debuggers

1 The term *tool* widely refers to programs executed by the user or system administrator on a
2 command line. Tools frequently interact with either the SMS, user applications, or both to perform
3 administrative and support functions. For example, a debugger tool might be used to remotely
4 control the processes of a parallel application, monitoring their behavior on a step-by-step basis.
5 Historically, such tools were custom-written for each specific host environment due to the
6 customized and/or proprietary nature of the environment's interfaces.

7 The advent of PMIx offers the possibility for creating portable tools capable of interacting with
8 multiple RMs without modification. Possible use-cases include:

- 9 • querying the status of scheduling queues and estimated allocation time for various resource
10 options
- 11 • job submission and allocation requests
- 12 • querying job status for executing applications
- 13 • launching, monitoring, and debugging applications

14 Enabling these capabilities requires some extensions to the PMIx Standard (both in terms of APIs
15 and attributes), and utilization of client-side APIs for more tool-oriented purposes.

16 This chapter defines specific APIs related to tools, provides tool developers with an overview of the
17 support provided by PMIx, and serves to guide RM vendors regarding roles and responsibilities of
18 RMs to support tools. As the number of tool-specific APIs and attributes is fairly small, the bulk of
19 the chapter serves to provide a "theory of operation" for tools and debuggers. Description of the
20 APIs themselves is therefore deferred to the Section 18.5 later in the chapter.

21 18.1 Connection Mechanisms

22 The key to supporting tools lies in providing mechanisms by which a tool can connect to a PMIx
23 server. Application processes are able to connect because their local RM daemon provides them
24 with the necessary contact information upon execution. A command-line tool, however, isn't
25 spawned by an RM daemon, and therefore lacks the information required for rendezvous with a
26 PMIx server.

27 Once a tool has started, it initializes PMIx as a tool (via `PMIx_tool_init`) if its access is
28 restricted to PMIx-based informational services such as `PMIx_Query_info`. However, if the

1 tool intends to start jobs, then it must include the **PMIX_LAUNCHER** attribute to inform the library
2 of that intent so that the library can initialize and provide access to the corresponding support.

3 Support for tools requires that the PMIx server be initialized with an appropriate attribute
4 indicating that tool connections are to be allowed. Separate attributes are provided to "fine-tune"
5 this permission by allowing the environment to independently enable (or disable) connections from
6 tools executing on nodes other than the one hosting the server itself. The PMIx server library shall
7 provide an opportunity for the host environment to authenticate and approve each connection
8 request from a specific tool by calling the **pmix_server_tool_connection_fn_t** "hook"
9 provided in the server module for that purpose. Servers in environments that do not provide this
10 "hook" shall automatically reject all tool connection requests.

11 Tools can connect to any local or remote PMIx server provided they are either explicitly given the
12 required connection information, or are able to discover it via one of several defined rendezvous
13 protocols. Connection discovery centers around the existence of *rendezvous files* containing the
14 necessary connection information, as illustrated in Fig. 18.1.




Figure 18.1.: Tool rendezvous files

15 The contents of each rendezvous file are specific to a given PMIx implementation, but should at
16 least contain the namespace and rank of the server along with its connection URI. Note that tools
17 linked to one PMIx implementation are therefore unlikely to successfully connect to PMIx server
18 libraries from another implementation.

19 The top of the directory tree is defined by either the **PMIX_SYSTEM_TMPDIR** attribute (if given)
20 or the **TMPDIR** environmental variable. PMIx servers that are designated as *system servers* by
21 including the **PMIX_SERVER_SYSTEM_SUPPORT** attribute when calling
22 **PMIx_server_init** will create a rendezvous file in this top-level directory. The filename will
23 be of the form *pmix.sys.hostname*, where *hostname* is the string returned by the **gethostname**
24 system call. Note that only one PMIx server on a node can be designated as the system server.

25 Non-system PMIx servers will create a set of three rendezvous files in the directory defined by
26 either the **PMIX_SERVER_TMPDIR** attribute or the **TMPDIR** environmental variable:

- 1 • *pmix.host.tool.namespace* where *host* is the string returned by the `gethostname` system call and
2 *namespace* is the namespace of the server.
- 3 • *pmix.host.tool.pid* where *host* is the string returned by the `gethostname` system call and *pid* is
4 the PID of the server.
- 5 • *pmix.host.tool* where *host* is the string returned by the `gethostname` system call. Note that
6 servers which are not given a namespace-specific `PMIX_SERVER_TMPDIR` attribute may not
7 generate this file due to conflicts should multiple servers be present on the node.

8 The files are identical and may be implemented as symlinks to a single instance. The individual file
9 names are composed so as to aid the search process should a tool wish to connect to a server
10 identified by its namespace or PID.

11 Servers will additionally provide a rendezvous file in any given location if the path (either absolute
12 or relative) and filename is specified either during `PMIx_server_init` using the
13 `PMIX_LAUNCHER_RENDEZVOUS_FILE` attribute, or by the `PMIX_LAUNCHER_RNDZ_FILE`
14 environmental variable prior to executing the process containing the server. This latter mechanism
15 may be the preferred mechanism for tools such as debuggers that need to fork/exec a launcher (e.g.,
16 "mpiexec") and then rendezvous with it. This is described in more detail in Section 18.2.2.

17 Rendezvous file ownerships are set to the UID and GID of the server that created them, with
18 permissions set according to the desires of the implementation and/or system administrator policy.
19 All connection attempts are first governed by read access privileges to the target rendezvous file -
20 thus, the combination of permissions, UID, and GID of the rendezvous files act as a first-level of
21 security for tool access.

22 A tool may connect to as many servers at one time as the implementation supports, but is limited to
23 designating only one such connection as its *primary* server. This is done to avoid confusion when
24 the tool calls an API as to which server should service the request. The first server the tool connects
25 to is automatically designated as the *primary* server.

26 Tools are allowed to change their primary server at any time via the `PMIx_tool_set_server`
27 API, and to connect/disconnect from a server as many times as desired. Note that standing requests
28 (e.g., event registrations) with the current primary server may be lost and/or may not be transferred
29 when transitioning to another primary server - PMIx implementors are not required to maintain or
30 transfer state across tool-server connections.

31 Tool process identifiers are assigned by one of the following methods:

- 32 • If `PMIX_TOOL_NAMESPACE` is given, then the namespace of the tool will be assigned that value.
33 – If `PMIX_TOOL_RANK` is also given, then the rank of the tool will be assigned that value.
34 – If `PMIX_TOOL_RANK` is not given, then the rank will be set to a default value of zero.
- 35 • If a process ID is not provided and the tool connects to a server, then one will be assigned by the
36 host environment upon connection to that server.

- If a process ID is not provided and the tool does not connect to a server (e.g., if `PMIX_TOOL_DO_NOT_CONNECT` is given), then the tool shall self-assign a unique identifier. This is often done using some combination involving hostname and PID.

Tool process identifiers remain constant across servers. Thus, it is critical that a system-wide unique namespace be provided if the tool itself sets the identifier, and that host environments provide a system-wide unique identifier in the case where the identifier is set by the server upon connection. The host environment is required to reject any connection request that fails to meet this criterion.

For simplicity, the following descriptions will refer to the:

- `PMIX_SYSTEM_TMPDIR` as the directory specified by either the `PMIX_SYSTEM_TMPDIR` attribute (if given) or the `TMPDIR` environmental variable.
- `PMIX_SERVER_TMPDIR` as the directory specified by either the `PMIX_SERVER_TMPDIR` attribute or the `TMPDIR` environmental variable.

The rendezvous methods are automatically employed for the initial tool connection during `PMIx_tool_init` unless the `PMIX_TOOL_DO_NOT_CONNECT` attribute is specified, and on all subsequent calls to `PMIx_tool_attach_to_server`.

18.1.1 Rendezvousing with a local server

Connection to a local PMIx server is pursued according to the following precedence chain based on attributes contained in the call to the `PMIx_tool_init` or `PMIx_tool_attach_to_server` APIs. Servers to which the tool already holds a connection will be ignored. Except where noted, the PMIx library will return an error if the specified file cannot be found, the caller lacks permissions to read it, or the server specified within the file does not respond to or accept the connection — the library will not proceed to check for other connection options as the user specified a particular one to use.

Note that the PMIx implementation may choose to introduce a "delayed connection" protocol between steps in the precedence chain - i.e., the library may cycle several times, checking for creation of the rendezvous file each time after a delay of some period of time, thereby allowing the tool to wait for the server to create the rendezvous file before either returning an error or continuing to the next step in the chain.

- If `PMIX_TOOL_ATTACHMENT_FILE` is given, then the tool will attempt to read the specified file and connect to the server based on the information contained within it. The format of the attachment file is identical to the rendezvous files described in earlier in this section. An error will be returned if the specified file cannot be found.
- If `PMIX_SERVER_URI` or `PMIX_TCP_URI` is given, then connection will be attempted to the server at the specified URI. Note that it is an error for both of these attributes to be specified. `PMIX_SERVER_URI` is the preferred method as it is more generalized — `PMIX_TCP_URI` is provided for those cases where the user specifically wants to use a TCP transport for the connection and wants to error out if one isn't available or cannot be used.

- 1 • If `PMIX_SERVER_PIDINFO` was provided, then the tool will search for a rendezvous file
2 created by a PMIx server of the given PID in the `PMIX_SERVER_TMPDIR` directory. An error
3 will be returned if a matching rendezvous file cannot be found.
- 4 • If `PMIX_SERVER_NAMESPACE` is given, then the tool will search for a rendezvous file created by a
5 PMIx server of the given namespace in the `PMIX_SERVER_TMPDIR` directory. An error will
6 be returned if a matching rendezvous file cannot be found.
- 7 • If `PMIX_CONNECT_TO_SYSTEM` is given, then the tool will search for a system-level
8 rendezvous file created by a PMIx server in the `PMIX_SYSTEM_TMPDIR` directory. An error
9 will be returned if a matching rendezvous file cannot be found.
- 10 • If `PMIX_CONNECT_SYSTEM_FIRST` is given, then the tool will look for a system-level
11 rendezvous file created by a PMIx server in the `PMIX_SYSTEM_TMPDIR` directory. If found,
12 then the tool will attempt to connect to it. In this case, no error will be returned if the rendezvous
13 file is not found or connection is refused — the PMIx library will silently continue to the next
14 option.
- 15 • By default, the tool will search the directory tree under the `PMIX_SERVER_TMPDIR` directory
16 for rendezvous files of PMIx servers, attempting to connect to each it finds until one accepts the
17 connection. If no rendezvous files are found, or all contacted servers refuse connection, then the
18 PMIx library will return an error. No "delayed connection" protocols may be utilized at this point.

19 Note that there can be multiple local servers - one from the system plus others from launchers and
20 active jobs. The PMIx tool connection search method is not guaranteed to pick a particular server
21 unless directed to do so. Tools can obtain a list of servers available on their local node using the
22 `PMIx_Query_info` APIs with the `PMIX_QUERY_AVAIL_SERVERS` key.

23 18.1.2 Connecting to a remote server

24 Connecting to remote servers is complicated due to the lack of access to the previously-described
25 rendezvous files. Two methods are required to be supported, both based on the caller having explicit
26 knowledge of either connection information or a path to a local file that contains such information:

- 27 • If `PMIX_TOOL_ATTACHMENT_FILE` is given, then the tool will attempt to read the specified
28 file and connect to the server based on the information contained within it. The format of the
29 attachment file is identical to the rendezvous files described in earlier in this section.
- 30 • If `PMIX_SERVER_URI` or `PMIX_TCP_URI` is given, then connection will be attempted to the
31 server at the specified URI. Note that it is an error for both of these attributes to be specified.
32 `PMIX_SERVER_URI` is the preferred method as it is more generalized — `PMIX_TCP_URI` is
33 provided for those cases where the user specifically wants to use the TCP transport for the
34 connection and wants to error out if it isn't available or cannot be used.

35 Additional methods may be provided by particular PMIx implementations. For example, the tool
36 may use `ssh` to launch a *probe* process onto the remote node so that the probe can search the
37 `PMIX_SYSTEM_TMPDIR` and `PMIX_SERVER_TMPDIR` directories for rendezvous files,

1 relaying the discovered information back to the requesting tool. If sufficient information is found to
2 allow for remote connection, then the tool can use it to establish the connection. Note that this
3 method is not required to be supported - it is provided here as an example and left to the discretion
4 of PMIx implementors.

5 **18.1.3 Attaching to running jobs**

6 When attaching to a running job, the tool must connect to a PMIx server that is associated with that
7 job - e.g., a server residing in the host environment's local daemon that spawned one or more of the
8 job's processes, or the server residing in the launcher that is overseeing the job. Identifying an
9 appropriate server can sometimes prove challenging, particularly in an environment where multiple
10 job launchers may be in operation, possibly under control of the same user.

11 In cases where the user has only the one job of interest in operation on the local node (e.g., when
12 engaged in an interactive session on the node from which the launcher was executed), the normal
13 rendezvous file discovery method can often be used to successfully connect to the target job, even
14 in the presence of jobs executed by other users. The permissions and security authorizations can, in
15 many cases, reliably ensure that only the one connection can be made. However, this is not
16 guaranteed in all cases.

17 The most common method, therefore, for attaching to a running job is to specify either the PID of
18 the job's launcher or the namespace of the launcher's job (note that the launcher's namespace
19 frequently differs from the namespace of the job it has launched). Unless the application processes
20 themselves act as PMIx servers, connection must be to the servers in the daemons that oversee the
21 application. This is typically either daemons specifically started by the job's launcher process, or
22 daemons belonging to the host environment, that are responsible for starting the application's
23 processes and oversee their execution.

24 Identifying the correct PID or namespace can be accomplished in a variety of ways, including:

- 25 • Using typical OS or host environment tools to obtain a listing of active jobs and perusing those to
26 find the target launcher.
- 27 • Using a PMIx-based tool attached to a system-level server to query the active jobs and their
28 command lines, thereby identifying the application of interest and its associated launcher.
- 29 • Manually recording the PID of the launcher upon starting the job.

30 Once the namespace and/or PID of the target server has been identified, either of the previous
31 methods can be used to connect to it.

32 **18.1.4 Tool initialization attributes**

33 The following attributes are passed to the `PMIx_tool_init` API for use when initializing the
34 PMIx library.

35 `PMIX_TOOL_NAMESPACE` "pmix.tool.namespace" (char*)

1 Name of the namespace to use for this tool.
2 **PMIX_TOOL_RANK** "pmix.tool.rank" (uint32_t)
3 Rank of this tool.
4 **PMIX_LAUNCHER** "pmix.tool.launcher" (bool)
5 Tool is a launcher and needs to create rendezvous files.

6 18.1.5 Tool initialization environmental variables

7 The following environmental variables are used during **PMIx_tool_init** and
8 **PMIx_server_init** to control various rendezvous-related operations when the process is
9 started manually (e.g., on a command line) or by a fork/exec-like operation.

10 **PMIX_LAUNCHER_RNDZ_URI**

11 The spawned tool is to be connected back to the spawning tool using the given URI so that
12 the spawning tool can provide directives (e.g., a **PMIx_Spawn** command) to it.

13 **PMIX_LAUNCHER_RNDZ_FILE**

14 If the specified file does not exist, this variable contains the absolute path of the file where
15 the spawned tool is to store its connection information so that the spawning tool can connect
16 to it. If the file does exist, it contains the information specifying the server to which the
17 spawned tool is to connect.

18 **PMIX_KEEPA_LIVE_PIPE**

19 An integer **read**-end of a POSIX pipe that the tool should monitor for closure, thereby
20 indicating that the parent tool has terminated. Used, for example, when a tool fork/exec's an
21 intermediate launcher that should self-terminate if the originating tool exits.

22 Note that these environmental variables should be cleared from the environment after use and prior
23 to forking child processes to avoid potentially unexpected behavior by the child processes.

24 18.1.6 Tool connection attributes

25 These attributes are defined to assist PMIx-enabled tools to connect with a PMIx server by passing
26 them into either the **PMIx_tool_init** or the **PMIx_tool_attach_to_server** APIs - thus,
27 they are not typically accessed via the **PMIx_Get** API.

28 **PMIX_SERVER_PIDINFO** "pmix.srvr.pidinfo" (pid_t)

29 PID of the target PMIx server for a tool.

30 **PMIX_CONNECT_TO_SYSTEM** "pmix.cnct.sys" (bool)

31 The requester requires that a connection be made only to a local, system-level PMIx server.

32 **PMIX_CONNECT_SYSTEM_FIRST** "pmix.cnct.sys.first" (bool)

33 Preferentially, look for a system-level PMIx server first.

34 **PMIX_SERVER_URI** "pmix.srvr.uri" (char*)

35 URI of the PMIx server to be contacted.

36 **PMIX_SERVER_HOSTNAME** "pmix.srvr.host" (char*)

37 Host where target PMIx server is located.

38 **PMIX_CONNECT_MAX_RETRIES** "pmix.tool.mretries" (uint32_t)

1 Maximum number of times to try to connect to PMIx server - the default value is
2 implementation specific.

3 **PMIX_CONNECT_RETRY_DELAY** "pmix.tool.retry" (uint32_t)

4 Time in seconds between connection attempts to a PMIx server - the default value is
5 implementation specific.

6 **PMIX_TOOL_DO_NOT_CONNECT** "pmix.tool.nocon" (bool)

7 The tool wants to use internal PMIx support, but does not want to connect to a PMIx server.

8 **PMIX_TOOL_CONNECT_OPTIONAL** "pmix.tool.conopt" (bool)

9 The tool shall connect to a server if available, but otherwise continue to operate unconnected.

10 **PMIX_TOOL_ATTACHMENT_FILE** "pmix.tool.attach" (char*)

11 Pathname of file containing connection information to be used for attaching to a specific
12 server.

13 **PMIX_LAUNCHER_RENDEZVOUS_FILE** "pmix.tool.lncrnd" (char*)

14 Pathname of file where the launcher is to store its connection information so that the
15 spawning tool can connect to it.

16 **PMIX_PRIMARY_SERVER** "pmix.pri.srvr" (bool)

17 The server to which the tool is connecting shall be designated the *primary* server once
18 connection has been accomplished.

19 **PMIX_WAIT_FOR_CONNECTION** "pmix.wait.conn" (bool)

20 Wait until the specified process has connected to the requesting tool or server, or the
21 operation times out (if the **PMIX_TIMEOUT** directive is included in the request).

22 18.2 Launching Applications with Tools

23 Tool-directed launches require that the tool include the **PMIX_LAUNCHER** attribute when calling
24 **PMIx_tool_init**. Two launch modes are supported:

- 25 • *Direct launch* where the tool itself is directly responsible for launching all processes, including
26 debugger daemons, using either the RM or daemons launched by the tool – i.e., there is no
27 *intermediate launcher* (IL) such as *mpiexec*. The case where the tool is self-contained (i.e., uses
28 its own daemons without interacting with an external entity such as the RM) lies outside the
29 scope of this Standard; and
- 30 • *Indirect launch* where all processes are started via an IL such as *mpiexec* and the tool itself is not
31 directly involved in launching application processes or debugger daemons. Note that the IL may
32 utilize the RM to launch processes and/or daemons under the tool's direction.

33 Either of these methods can be executed interactively or by a batch script. Note that not all host
34 environments may support the direct launch method.

35 18.2.1 Direct launch

36 In the direct-launch use-case (Fig. 18.2), the tool itself performs the role of the launcher. Once
37 invoked, the tool connects to an appropriate PMIx server - e.g., a system-level server hosted by the

1 RM. The tool is responsible for assembling the description of the application to be launched (e.g.,
 2 by parsing its command line) into a spawn request containing an array of `pmix_app_t`
 3 applications and `pmix_info_t` job-level information. An allocation of resources may or may not
 4 have been made in advance – if not, then the spawn request must include allocation request
 5 information.




Figure 18.2.: Direct Launch

6 In addition to the attributes described in `PMIx_Spawn`, the tool may optionally wish to include the
 7 following tool-specific attributes in the `job_info` argument to that API (the debugger-related
 8 attributes are discussed in more detail in Section 18.4):

- 9 • `PMIX_FWD_STDIN` "`pmix.fwd.stdin`" (`pmix_rank_t`)
 10 The requester intends to push information from its `stdin` to the indicated process. The
 11 local spawn agent should, therefore, ensure that the `stdin` channel to that process
 12 remains available. A rank of `PMIX_RANK_WILDCARD` indicates that all processes in the
 13 spawned job are potential recipients. The requester will issue a call to `PMIx_IOF_push`
 14 to initiate the actual forwarding of information to specified targets - this attribute simply
 15 requests that the IL retain the ability to forward the information to the designated targets.
- 16 • `PMIX_FWD_STDOUT` "`pmix.fwd.stdout`" (`bool`)
 17 Requests that the ability to forward the `stdout` of the spawned processes be maintained.
 18 The requester will issue a call to `PMIx_IOF_pull` to specify the callback function and
 19 other options for delivery of the forwarded output.
- 20 • `PMIX_FWD_STDERR` "`pmix.fwd.stderr`" (`bool`)
 21 Requests that the ability to forward the `stderr` of the spawned processes be maintained.
 22 The requester will issue a call to `PMIx_IOF_pull` to specify the callback function and

1 other options for delivery of the forwarded output.

- 2 • **PMIX_FWD_STDDIAG** "pmix.fwd.stddiag" (bool)
3 Requests that the ability to forward the diagnostic channel (if it exists) of the spawned
4 processes be maintained. The requester will issue a call to **PMIx_IOF_pull** to specify
5 the callback function and other options for delivery of the forwarded output.
- 6 • **PMIX_IOF_CACHE_SIZE** "pmix.iof.csize" (uint32_t)
7 The requested size of the PMIx server cache in bytes for each specified channel. By
8 default, the server is allowed (but not required) to drop all bytes received beyond the max
9 size.
- 10 • **PMIX_IOF_DROP_OLDEST** "pmix.iof.old" (bool)
11 In an overflow situation, the PMIx server is to drop the oldest bytes to make room in the
12 cache.
- 13 • **PMIX_IOF_DROP_NEWEST** "pmix.iof.new" (bool)
14 In an overflow situation, the PMIx server is to drop any new bytes received until room
15 becomes available in the cache (default).
- 16 • **PMIX_IOF_BUFFERING_SIZE** "pmix.iof.bsize" (uint32_t)
17 Requests that IO on the specified channel(s) be aggregated in the PMIx tool library until
18 the specified number of bytes is collected to avoid being called every time a block of IO
19 arrives. The PMIx tool library will execute the callback and reset the collection counter
20 whenever the specified number of bytes becomes available. Any remaining buffered data
21 will be *flushed* to the callback upon a call to deregister the respective channel.
- 22 • **PMIX_IOF_BUFFERING_TIME** "pmix.iof.btime" (uint32_t)
23 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
24 size, this prevents IO from being held indefinitely while waiting for another payload to
25 arrive.
- 26 • **PMIX_IOF_OUTPUT_RAW** "pmix.iof.raw" (bool)
27 Do not buffer output to be written as complete lines - output characters as the stream
28 delivers them
- 29 • **PMIX_IOF_TAG_OUTPUT** "pmix.iof.tag" (bool)
30 Requests that output be prefixed with the nspace,rank of the source and a string
31 identifying the channel (**stdout**, **stderr**, etc.).
- 32 • **PMIX_IOF_TIMESTAMP_OUTPUT** "pmix.iof.ts" (bool)
33 Requests that output be marked with the time at which the data was received by the tool -
34 note that this will differ from the time at which the data was collected from the source.
- 35 • **PMIX_IOF_XML_OUTPUT** "pmix.iof.xml" (bool)
36 Requests that output be formatted in XML.
- 37 • **PMIX_IOF_RANK_OUTPUT** "pmix.iof.rank" (bool)
38 Tag output with the rank it came from

- 1 • **PMIX_IOF_OUTPUT_TO_FILE** "pmix.iof.file" (char*)
2 Direct application output into files of form "<filename>.<nospace>.<rank>.stdout" (for
3 **stdout**) and "<filename>.<nospace>.<rank>.stderr" (for **stderr**). If
4 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
5 created and both streams will be written into it.
- 6 • **PMIX_IOF_OUTPUT_TO_DIRECTORY** "pmix.iof.dir" (char*)
7 Direct application output into files of form "<directory>/<nospace>/rank.<rank>/stdout"
8 (for **stdout**) and "<directory>/<nospace>/rank.<rank>/stderr" (for **stderr**). If
9 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
10 created and both streams will be written into it.
- 11 • **PMIX_IOF_FILE_PATTERN** "pmix.iof.fpt" (bool)
12 Specified output file is to be treated as a pattern and not automatically annotated by
13 nspace, rank, or other parameters. The pattern can use %n for the namespace, and %r for
14 the rank wherever those quantities are to be placed. The resulting filename will be
15 appended with ".stdout" for the **stdout** stream and ".stderr" for the **stderr** stream. If
16 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
17 created and both streams will be written into it.
- 18 • **PMIX_IOF_FILE_ONLY** "pmix.iof.fonly" (bool)
19 Output only into designated files - do not also output a copy to the console's stdout/stderr
- 20 • **PMIX_IOF_MERGE_STDERR_STDOUT** "pmix.iof.mrg" (bool)
21 Merge stdout and stderr streams from application proc
- 22 • **PMIX_NOHUP** "pmix.nohup" (bool)
23 Any processes started on behalf of the calling tool (or the specified namespace, if such
24 specification is included in the list of attributes) should continue after the tool disconnects
25 from its server.
- 26 • **PMIX_NOTIFY_JOB_EVENTS** "pmix.note.jev" (bool)
27 Requests that the launcher generate the **PMIX_EVENT_JOB_START**,
28 **PMIX_LAUNCH_COMPLETE**, and **PMIX_EVENT_JOB_END** events. Each event is to
29 include at least the namespace of the corresponding job and a
30 **PMIX_EVENT_TIMESTAMP** indicating the time the event occurred. Note that the
31 requester must register for these individual events, or capture and process them by
32 registering a default event handler instead of individual handlers and then process the
33 events based on the returned status code. Another common method is to register one event
34 handler for all job-related events, with a separate handler for non-job events - see
35 **PMIx_Register_event_handler** for details.
- 36 • **PMIX_NOTIFY_COMPLETION** "pmix.notecomp" (bool)
37 Requests that the launcher generate the **PMIX_EVENT_JOB_END** event for normal or
38 abnormal termination of the spawned job. The event shall include the returned status code
39 (**PMIX_JOB_TERM_STATUS**) for the corresponding job; the identity (**PMIX_PROCID**)
40 and exit status (**PMIX_EXIT_CODE**) of the first failed process, if applicable; and a

1 **PMIX_EVENT_TIMESTAMP** indicating the time the termination occurred. Note that the
2 requester must register for the event or capture and process it within a default event
3 handler.

- 4 • **PMIX_LOG_JOB_EVENTS** "pmix.log.jev" (bool)
5 Requests that the launcher log the **PMIX_EVENT_JOB_START**,
6 **PMIX_LAUNCH_COMPLETE**, and **PMIX_EVENT_JOB_END** events using **PMIx_Log**,
7 subject to the logging attributes of Section 13.4.3.
- 8 • **PMIX_LOG_COMPLETION** "pmix.logcomp" (bool)
9 Requests that the launcher log the **PMIX_EVENT_JOB_END** event for normal or
10 abnormal termination of the spawned job using **PMIx_Log**, subject to the logging
11 attributes of Section 13.4.3. The event shall include the returned status code
12 (**PMIX_JOB_TERM_STATUS**) for the corresponding job; the identity (**PMIX_PROCID**)
13 and exit status (**PMIX_EXIT_CODE**) of the first failed process, if applicable; and a
14 **PMIX_EVENT_TIMESTAMP** indicating the time the termination occurred.
- 15 • **PMIX_DEBUG_STOP_ON_EXEC** "pmix.dbg.exec" (bool)
16 Included in either the **pmix_info_t** array in a **pmix_app_t** description (if the
17 directive applies only to that application) or in the *job_info* array if it applies to all
18 applications in the given spawn request. Indicates that the application is being spawned
19 under a debugger, and that the local launch agent is to pause the resulting application
20 processes on first instruction for debugger attach. The launcher (RM or IL) is to generate
21 the **PMIX_LAUNCH_COMPLETE** event when all processes are stopped at the exec point.
- 22 • **PMIX_DEBUG_STOP_IN_INIT** "pmix.dbg.init" (bool)
23 Included in either the **pmix_info_t** array in a **pmix_app_t** description (if the
24 directive applies only to that application) or in the *job_info* array if it applies to all
25 applications in the given spawn request. Indicates that the specified application is being
26 spawned under a debugger. The PMIx client library in each resulting application process
27 shall notify its PMIx server that it is pausing and then pause during **PMIx_Init** of the
28 spawned processes until either released by debugger modification of an appropriate
29 variable or receipt of the **PMIX_DEBUGGER_RELEASE** event. The launcher (RM or IL)
30 is responsible for generating the **PMIX_READY_FOR_DEBUG** event (stipulating a
31 breakpoint of *pmix-init*) when all processes have reached the pause point.
- 32 • **PMIX_DEBUG_STOP_IN_APP** "pmix.dbg.notify" (varies)
33 Direct specified ranks to stop at application-specific point and notify they are
34 ready-to-debug. The attribute's value can be any of three data types:
35 – bool - true indicating all ranks
36 – **pmix_rank_t** - the rank of one proc, or **PMIX_RANK_WILDCARD** for all
37 – a **pmix_data_array_t** if an array of individual processes are specified
38 The resulting application processes are to notify their server (by generating the
39 **PMIX_READY_FOR_DEBUG** event) when they reach some application-determined
40 location - the event shall include the **PMIX_BREAKPOINT** attribute indicating where the

1 application has stopped. The application shall pause at that point until released by
2 debugger modification of an appropriate variable. The launcher (RM or IL) is responsible
3 for generating the **PMIX_READY_FOR_DEBUG** event when all processes have indicated
4 they are at the pause point.

Advice to users

5 The **PMIX_IOF_FILE_ONLY** indicates output is directed to files and no copy is sent back to the
6 application. For example, this can be combined with **PMIX_IOF_OUTPUT_TO_FILE** or
7 **PMIX_IOF_OUTPUT_TO_DIRECTORY** to only output to files.

8 The tool then calls the **PMIx_Spawn** API so that the PMIx library can communicate the spawn
9 request to the server.

10 Upon receipt, the PMIx server library passes the spawn request to its host RM daemon for
11 processing via the **pmix_server_spawn_fn_t** server module function. If this callback was not
12 provided, then the PMIx server library will return the **PMIX_ERR_NOT_SUPPORTED** error status.

13 If an allocation must be made, then the host environment is responsible for communicating the
14 request to its associated scheduler. Once resources are available, the host environment initiates the
15 launch process to start the job. The host environment must parse the spawn request for relevant
16 directives, returning an error if any required directive cannot be supported. Optional directives may
17 be ignored if they cannot be supported.

18 Any error while executing the spawn request must be returned by **PMIx_Spawn** to the requester.
19 Once the spawn request has succeeded in starting the specified processes, the request will return
20 **PMIX_SUCCESS** back to the requester along with the namespace of the started job. Upon
21 termination of the spawned job, the host environment must generate a **PMIX_EVENT_JOB_END**
22 event for normal or abnormal termination if requested to do so. The event shall include:

- 23 • the returned status code (**PMIX_JOB_TERM_STATUS**) for the corresponding job;
- 24 • the identity (**PMIX_PROCID**) and exit status (**PMIX_EXIT_CODE**) of the first failed process, if
25 applicable;
- 26 • a **PMIX_EVENT_TIMESTAMP** indicating the time the termination occurred; plus
- 27 • any other info provided by the host environment.

1 18.2.2 Indirect launch

2 In the indirect launch use-case, the application processes are started via an intermediate launcher
3 (e.g., *mpiexec*) that is itself started by the tool (see Fig 18.3). Thus, at a high level, this is a
4 two-stage launch procedure to start the application: the tool (henceforth referred to as the *initiator*)
5 starts the IL, which then starts the applications. In practice, additional steps may be involved if, for
6 example, the IL starts its own daemons to shepherd the application processes.

7 A key aspect of this operational mode is the avoidance of any requirement that the initiator parse
8 and/or understand the command line of the IL. Instead, the indirect launch procedure supports
9 either of two methods: one where the initiator assumes responsibility for parsing its command line
10 to obtain the application as well as the IL and its options, and another where the initiator defers the
11 command line parsing to the IL. Both of these methods are described in the following sections.

12 18.2.2.1 Initiator-based command line parsing

13 This method utilizes a first call to the **PMIx_Spawn** API to start the IL itself, and then uses a
14 second call to **PMIx_Spawn** to request that the IL spawn the actual job. The burden of analyzing
15 the initial command line to separately identify the IL's command line from the application itself
16 falls upon the initiator. An example is provided below:

```
17 $ initiator --launcher "mpiexec --verbose" -n 3 ./app <appoptions>
```

18 The initiator spawns the IL using the same procedure for launching an application - it begins by
19 assembling the description of the IL into a spawn request containing an array of **pmix_app_t** and
20 **pmix_info_t** job-level information. Note that this step does not include any information
21 regarding the application itself - only the launcher is included. In addition, the initiator must
22 include the rendezvous URI in the environment so the IL knows how to connect back to it.

23 An allocation of resources for the IL itself may or may not be required – if it is, then the allocation
24 must be made in advance or the spawn request must include allocation request information.

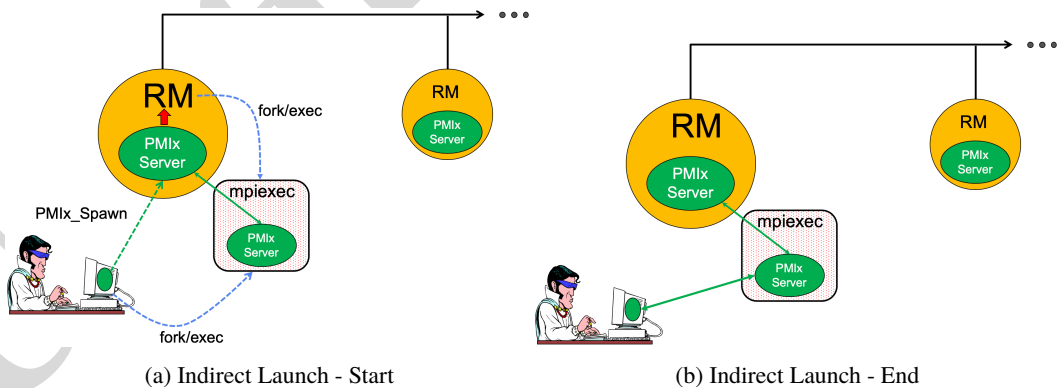


Figure 18.3.: Indirect launch procedure

1 The initiator may optionally wish to include the following tool-specific attributes in the *job_info*
2 argument to **PMIx_Spawn** - note that these attributes refer only to the behavior of the IL itself and
3 not the eventual job to be launched:

- 4 • **PMIX_FWD_STDIN** "pmix.fwd.stdin" (**pmix_rank_t**)
5 The requester intends to push information from its **stdin** to the indicated process. The
6 local spawn agent should, therefore, ensure that the **stdin** channel to that process
7 remains available. A rank of **PMIX_RANK_WILDCARD** indicates that all processes in the
8 spawned job are potential recipients. The requester will issue a call to **PMIx_IOF_push**
9 to initiate the actual forwarding of information to specified targets - this attribute simply
10 requests that the IL retain the ability to forward the information to the designated targets.
- 11 • **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (**bool**)
12 Requests that the ability to forward the **stdout** of the spawned processes be maintained.
13 The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and
14 other options for delivery of the forwarded output.
- 15 • **PMIX_FWD_STDERR** "pmix.fwd.stderr" (**bool**)
16 Requests that the ability to forward the **stderr** of the spawned processes be maintained.
17 The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and
18 other options for delivery of the forwarded output.
- 19 • **PMIX_FWD_STDDIAG** "pmix.fwd.stddiag" (**bool**)
20 Requests that the ability to forward the diagnostic channel (if it exists) of the spawned
21 processes be maintained. The requester will issue a call to **PMIx_IOF_pull** to specify
22 the callback function and other options for delivery of the forwarded output.
- 23 • **PMIX_IOF_CACHE_SIZE** "pmix.iof.csize" (**uint32_t**)
24 The requested size of the PMIx server cache in bytes for each specified channel. By
25 default, the server is allowed (but not required) to drop all bytes received beyond the max
26 size.
- 27 • **PMIX_IOF_DROP_OLDEST** "pmix.iof.old" (**bool**)
28 In an overflow situation, the PMIx server is to drop the oldest bytes to make room in the
29 cache.
- 30 • **PMIX_IOF_DROP_NEWEST** "pmix.iof.new" (**bool**)
31 In an overflow situation, the PMIx server is to drop any new bytes received until room
32 becomes available in the cache (default).
- 33 • **PMIX_IOF_BUFFERING_SIZE** "pmix.iof.bsize" (**uint32_t**)
34 Requests that IO on the specified channel(s) be aggregated in the PMIx tool library until
35 the specified number of bytes is collected to avoid being called every time a block of IO
36 arrives. The PMIx tool library will execute the callback and reset the collection counter
37 whenever the specified number of bytes becomes available. Any remaining buffered data
38 will be *flushed* to the callback upon a call to deregister the respective channel.
- 39 • **PMIX_IOF_BUFFERING_TIME** "pmix.iof.btime" (**uint32_t**)

1 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
2 size, this prevents IO from being held indefinitely while waiting for another payload to
3 arrive.

- 4 • **PMIX_IOF_TAG_OUTPUT** "pmix.iof.tag" (bool)
5 Requests that output be prefixed with the nspace,rank of the source and a string
6 identifying the channel (**stdout**, **stderr**, etc.).
- 7 • **PMIX_IOF_TIMESTAMP_OUTPUT** "pmix.iof.ts" (bool)
8 Requests that output be marked with the time at which the data was received by the tool -
9 note that this will differ from the time at which the data was collected from the source.
- 10 • **PMIX_IOF_XML_OUTPUT** "pmix.iof.xml" (bool)
11 Requests that output be formatted in XML.
- 12 • **PMIX_NOHUP** "pmix.nohup" (bool)
13 Any processes started on behalf of the calling tool (or the specified namespace, if such
14 specification is included in the list of attributes) should continue after the tool disconnects
15 from its server.
- 16 • **PMIX_LAUNCHER_DAEMON** "pmix.lnch.dmn" (char*)
17 Path to executable that is to be used as the backend daemon for the launcher. This replaces
18 the launcher's own daemon with the specified executable. Note that the user is therefore
19 responsible for ensuring compatibility of the specified executable and the host launcher.
- 20 • **PMIX_FORKEXEC_AGENT** "pmix.frkex.agnt" (char*)
21 Path to executable that the launcher's backend daemons are to fork/exec in place of the
22 actual application processes. The fork/exec agent shall connect back (as a PMIx tool) to
23 the launcher's daemon to receive its spawn instructions, and is responsible for starting the
24 actual application process it replaced. See Section 18.4.3 for details.
- 25 • **PMIX_EXEC_AGENT** "pmix.exec.agnt" (char*)
26 Path to executable that the launcher's backend daemons are to fork/exec in place of the
27 actual application processes. The launcher's daemon shall pass the full command line of
28 the application on the command line of the exec agent, which shall not connect back to the
29 launcher's daemon. The exec agent is responsible for exec'ing the specified application
30 process in its own place. See Section 18.4.3 for details.
- 31 • **PMIX_DEBUG_STOP_IN_INIT** "pmix.dbg.init" (bool)
32 Included in either the **pmix_info_t** array in a **pmix_app_t** description (if the
33 directive applies only to that application) or in the *job_info* array if it applies to all
34 applications in the given spawn request. Indicates that the specified application is being
35 spawned under a debugger. The PMIx client library in each resulting application process
36 shall notify its PMIx server that it is pausing and then pause during **PMIx_Init** of the
37 spawned processes until either released by debugger modification of an appropriate
38 variable or receipt of the **PMIX_DEBUGGER_RELEASE** event. The launcher (RM or IL)
39 is responsible for generating the **PMIX_READY_FOR_DEBUG** event (stipulating a

breakpoint of `pmix-init`) when all processes have reached the pause point. In this context, the initiator is directing the IL to stop in `PMIx_tool_init`. This gives the initiator a chance to connect to the IL and register for events prior to the IL launching the application job.

and the following optional variables in the environment of the IL:

- `PMIX_KEEPALIVE_PIPE` - an integer `read`-end of a POSIX pipe that the IL should monitor for closure, thereby indicating that the initiator has terminated.

The initiator then calls the `PMIx_Spawn` API so that the PMIx library can either communicate the spawn request to a server (if connected to one), or locally spawn the IL itself if not connected to a server and the PMIx implementation includes self-spawn support. `PMIx_Spawn` shall return an error if neither of these conditions is met.

When initialized by the IL, the `PMIx_tool_init` function must perform two operations:

- check for the presence of the `PMIX_KEEPALIVE_PIPE` environmental variable - if provided, then the library shall monitor the pipe for closure, providing a `PMIX_EVENT_JOB_END` event when the pipe closes (thereby indicating the termination of the initiator). The IL should register for this event after completing `PMIx_tool_init` - the initiator's namespace can be obtained via a call to `PMIx_Get` with the `PMIX_PARENT_ID` key. Note that this feature will only be available if the spawned IL is local to the initiator.
- check for the `PMIX_LAUNCHER_RNDZ_URI` environmental parameter - if found, the library shall connect back to the initiator using the `PMIx_tool_attach_to_server` API, retaining its current server as its primary server.

Once the IL completes `PMIx_tool_init`, it must register for the `PMIX_EVENT_JOB_END` termination event and then idle until receiving that event - either directly from the initiator, or from the PMIx library upon detecting closure of the keepalive pipe. The IL idles in the intervening time as it is solely acting as a relay (if connected to a server that is performing the actual application launch) or as a PMIx server responding to spawn requests.

Upon return from the `PMIx_Spawn` API, the initiator should set the spawned IL as its primary server using the `PMIx_tool_set_server` API with the namespace returned by `PMIx_Spawn` and any valid rank (a rank of zero would ordinarily be used as only one IL process is typically started). It is advisable to set a connection timeout value when calling this function. The initiator can then proceed to spawn the actual application according to the procedure described in Section 18.2.1.

18.2.2.2 IL-based command line parsing

In the case where the initiator cannot parse its command line, it must defer that parsing to the IL. A common example is provided below:

```
$ initiator mpiexec --verbose -n 3 ./app <appoptions>
```

For this situation, the initiator proceeds as above with only one notable exception: instead of calling `PMIx_Spawn` twice (once to start the IL and again to start the actual application), the initiator only calls that API one time:

- 1 • The *app* parameter passed to the spawn request contains only one `pmix_app_t` that contains
2 the entire command line, including both launcher and application(s).
- 3 • The launcher executable must be in the *app.cmd* field and in *app.argv[0]*, with the rest of the
4 command line appended to the *app.argv* array.
- 5 • Any job-level directives for the IL itself (e.g., `PMIX_FORKEXEC_AGENT` or
6 `PMIX_FWD_STDOUT`) are included in the *job_info* parameter of the call to `PMIx_Spawn`.
- 7 • The job-level directives must include both the `PMIX_SPAWN_TOOL` attribute indicating that the
8 initiator is spawning a tool, and the `PMIX_DEBUG_STOP_IN_INIT` attribute directing the IL
9 to stop during the call to `PMIx_tool_init`. The latter directive allows the initiator to connect
10 to the IL prior to launch of the application.
- 11 • The `PMIX_LAUNCHER_RNDZ_URI` and `PMIX_KEEPLIVE_PIPE` environmental variables
12 are provided to the launcher in its environment via the *app.env* field.
- 13 • The IL must use `PMIx_Get` with the `PMIX_LAUNCH_DIRECTIVES` key to obtain any
14 initiator-provided directives (e.g., `PMIX_DEBUG_STOP_IN_INIT` or
15 `PMIX_DEBUG_STOP_ON_EXEC`) aimed at the application(s) it will spawn.

16 Upon return from `PMIx_Spawn`, the initiator must:

- 17 • use the `PMIx_tool_set_server` API to set the spawned IL as its primary server
- 18 • register with that server to receive the `PMIX_LAUNCH_COMPLETE` event. This allows the
19 initiator to know when the IL has completed launch of the application
- 20 • release the IL from its "hold" in `PMIx_tool_init` by issuing the
21 `PMIX_DEBUGGER_RELEASE` event, specifying the IL as the custom range. Upon receipt of the
22 event, the IL is free to parse its command line, apply any provided directives, and execute the
23 application.

24 Upon receipt of the `PMIX_LAUNCH_COMPLETE` event, the initiator should register to receive
25 notification of completion of the returned namespace of the application. Receipt of the
26 `PMIX_EVENT_JOB_END` event provides a signal that the initiator may itself terminate.

27 18.2.3 Tool spawn-related attributes

28 Tools are free to utilize the spawn attributes available to applications (see 12.2.4) when
29 constructing a spawn request, but can also utilize the following attributes that are specific to
30 tool-based spawn operations:

31 `PMIX_FWD_STDIN "pmix.fwd.stdin" (pmix_rank_t)`

32 The requester intends to push information from its `stdin` to the indicated process. The
33 local spawn agent should, therefore, ensure that the `stdin` channel to that process remains
34 available. A rank of `PMIX_RANK_WILDCARD` indicates that all processes in the spawned
35 job are potential recipients. The requester will issue a call to `PMIx_IOF_push` to initiate
36 the actual forwarding of information to specified targets - this attribute simply requests that
37 the IL retain the ability to forward the information to the designated targets.

1 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)
 2 Requests that the ability to forward the `stdout` of the spawned processes be maintained.
 3 The requester will issue a call to `PMIx_IOF_pull` to specify the callback function and
 4 other options for delivery of the forwarded output.

5 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)
 6 Requests that the ability to forward the `stderr` of the spawned processes be maintained.
 7 The requester will issue a call to `PMIx_IOF_pull` to specify the callback function and
 8 other options for delivery of the forwarded output.

9 **PMIX_FWD_STDDIAG** "pmix.fwd.stddiag" (bool)
 10 Requests that the ability to forward the diagnostic channel (if it exists) of the spawned
 11 processes be maintained. The requester will issue a call to `PMIx_IOF_pull` to specify the
 12 callback function and other options for delivery of the forwarded output.

13 **PMIX_NOHUP** "pmix.nohup" (bool)
 14 Any processes started on behalf of the calling tool (or the specified namespace, if such
 15 specification is included in the list of attributes) should continue after the tool disconnects
 16 from its server.

17 **PMIX_LAUNCHER_DAEMON** "pmix.lnch.dmn" (char*)
 18 Path to executable that is to be used as the backend daemon for the launcher. This replaces
 19 the launcher's own daemon with the specified executable. Note that the user is therefore
 20 responsible for ensuring compatibility of the specified executable and the host launcher.

21 **PMIX_FORKEXEC_AGENT** "pmix.frkex.agnt" (char*)
 22 Path to executable that the launcher's backend daemons are to fork/exec in place of the actual
 23 application processes. The fork/exec agent shall connect back (as a `PMIx` tool) to the
 24 launcher's daemon to receive its spawn instructions, and is responsible for starting the actual
 25 application process it replaced. See Section 18.4.3 for details.

26 **PMIX_EXEC_AGENT** "pmix.exec.agnt" (char*)
 27 Path to executable that the launcher's backend daemons are to fork/exec in place of the actual
 28 application processes. The launcher's daemon shall pass the full command line of the
 29 application on the command line of the exec agent, which shall not connect back to the
 30 launcher's daemon. The exec agent is responsible for exec'ing the specified application
 31 process in its own place. See Section 18.4.3 for details.

32 **PMIX_LAUNCH_DIRECTIVES** "pmix.lnch.dirs" (pmix_data_array_t*)
 33 Array of `pmix_info_t` containing directives for the launcher - a convenience attribute for
 34 retrieving all directives with a single call to `PMIx_Get`.

35 18.2.4 Tool rendezvous-related events

36 The following constants refer to events relating to rendezvous of a tool and launcher during spawn
 37 of the IL.

38 **PMIX_LAUNCHER_READY** -155 An application launcher (e.g., `mpiexec`) shall generate this
 39 event to signal a tool that started it that the launcher is ready to receive directives/commands
 40 (e.g., `PMIx_Spawn`). This is only used when the initiator is able to parse the command line
 41 itself, or the launcher is started as a persistent Distributed Virtual Machine (DVM).

1 18.3 IO Forwarding

2 Underlying the operation of many tools is a common need to forward **stdin** from the tool to
3 targeted processes, and to return **stdout/stderr** from those processes to the tool (e.g., for
4 display on the user's console). Historically, each tool developer was responsible for creating their
5 own IO forwarding subsystem. However, the introduction of PMIx as a standard mechanism for
6 interacting between applications and the host environment has made it possible to relieve tool
7 developers of this burden.

8 This section defines functions by which tools can request forwarding of input/output to/from other
9 processes and serves as a design guide to:

- 10 • provide tool developers with an overview of the expected behavior of the PMIx IO forwarding
11 support;
- 12 • guide RM vendors regarding roles and responsibilities expected of the RM to support IO
13 forwarding; and
- 14 • provide insight into the thinking of the PMIx community behind the definition of the PMIx IO
15 forwarding APIs.

16 Note that the forwarding of IO via PMIx requires that both the host environment and the tool
17 support PMIx, but does not impose any similar requirements on the application itself.

18 The responsibility of the host environment in forwarding of IO falls into the following areas:

- 19 • Capturing output from specified processes.
- 20 • Forwarding that output to the host of the PMIx server library that requested it.
- 21 • Delivering that payload to the PMIx server library via the [PMIx_server_IOF_deliver](#) API
22 for final dispatch to the requesting tool.

23 It is the responsibility of the PMIx library to buffer, format, and deliver the payload to the
24 requesting client. This may require caching of output until a forwarding registration is received, as
25 governed by the corresponding IO forwarding attributes of Section 18.3.5 that are supported by the
26 implementation.

27 18.3.1 Forwarding stdout/stderr

28 At an appropriate point in its operation (usually during startup), a tool will utilize the
29 [PMIx_tool_init](#) function to connect to a PMIx server. The PMIx server can be hosted by an
30 RM daemon or could be embedded in a library-provided starter program such as *mpiexec* - in terms
31 of IO forwarding, the operations remain the same either way. For purposes of this discussion, we
32 will assume the server is in an RM daemon and that the application processes are directly launched
33 by the RM, as shown in Fig 18.4.

34 Once the tool has connected to the target server, it can request that processes be spawned on its
35 behalf or that output from a specified set of existing processes in a given executing application be




Figure 18.4.: Forwarding stdout/stderr

1 forwarded to it. Requests to spawn processes should include the `PMIX_FWD_STDIN`,
 2 `PMIX_FWD_STDOUT`, and/or `PMIX_FWD_STDERR` attributes if the tool intends to request that
 3 the corresponding streams be forwarded at some point during execution.

4 Note that requests to capture output from existing processes via the `PMIX_IOF_PULL` API, and/or
 5 to forward input to specified processes via the `PMIX_IOF_PUSH` API, can only succeed if the
 6 required attributes to retain that ability were passed when the corresponding job was spawned. The
 7 host is required to return an error for all such requests in cases where this condition is not met.

8 Two modes are supported when requesting that the host forward standard output/error via the
 9 `PMIX_IOF_PULL` API - these can be controlled by including one of the following attributes in the
 10 *info* array passed to that function:

- 11 • `PMIX_IOF_COPY` "`pmix.iof.cpy`" (`bool`)
 12 Requests that the host environment deliver a copy of the specified output stream(s) to the
 13 tool, letting the stream(s) continue to also be delivered to the default location. This allows
 14 the tool to tap into the output stream(s) without redirecting it from its current final
 15 destination.
- 16 • `PMIX_IOF_REDIRECT` "`pmix.iof.redir`" (`bool`)
 17 Requests that the host environment intercept the specified output stream(s) and deliver it
 18 to the requesting tool instead of its current final destination. This might be used, for
 19 example, during a debugging procedure to avoid injection of debugger-related output into
 20 the application's results file. The original output stream(s) destination is restored upon
 21 termination of the tool. This is the default mode of operation.

22 When requesting to forward `stdout/stderr`, the tool can specify several formatting options to

1 be used on the resulting output stream. These include:

- 2 • **PMIX_IOF_TAG_OUTPUT** "pmix.iof.tag" (bool)
3 Requests that output be prefixed with the nspace,rank of the source and a string
4 identifying the channel (**stdout**, **stderr**, etc.).
- 5 • **PMIX_IOF_TIMESTAMP_OUTPUT** "pmix.iof.ts" (bool)
6 Requests that output be marked with the time at which the data was received by the tool -
7 note that this will differ from the time at which the data was collected from the source.
- 8 • **PMIX_IOF_XML_OUTPUT** "pmix.iof.xml" (bool)
9 Requests that output be formatted in XML.
- 10 • **PMIX_IOF_RANK_OUTPUT** "pmix.iof.rank" (bool)
11 Tag output with the rank it came from
- 12 • **PMIX_IOF_OUTPUT_TO_FILE** "pmix.iof.file" (char*)
13 Direct application output into files of form "<filename>.<nspace>.<rank>.stdout" (for
14 **stdout**) and "<filename>.<nspace>.<rank>.stderr" (for **stderr**). If
15 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
16 created and both streams will be written into it.
- 17 • **PMIX_IOF_OUTPUT_TO_DIRECTORY** "pmix.iof.dir" (char*)
18 Direct application output into files of form "<directory>/<nspace>/rank.<rank>/stdout"
19 (for **stdout**) and "<directory>/<nspace>/rank.<rank>/stderr" (for **stderr**). If
20 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
21 created and both streams will be written into it.
- 22 • **PMIX_IOF_FILE_PATTERN** "pmix.iof.fpt" (bool)
23 Specified output file is to be treated as a pattern and not automatically annotated by
24 nspace, rank, or other parameters. The pattern can use %n for the namespace, and %r
25 for the rank wherever those quantities are to be placed. The resulting filename will be
26 appended with ".stdout" for the **stdout** stream and ".stderr" for the **stderr** stream. If
27 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
28 created and both streams will be written into it.
- 29 • **PMIX_IOF_FILE_ONLY** "pmix.iof.fonly" (bool)
30 Output only into designated files - do not also output a copy to the console's stdout/stderr
- 31 • **PMIX_IOF_MERGE_STDERR_STDOUT** "pmix.iof.mrg" (bool)
32 Merge stdout and stderr streams from application procs

33 The PMIx client in the tool is responsible for formatting the output stream. Note that output from
34 multiple processes will often be interleaved due to variations in arrival time - ordering of output is
35 not guaranteed across processes and/or nodes.

1 18.3.2 Forwarding stdin

2 A tool is not necessarily a child of the RM as it may have been started directly from the command
3 line. Thus, provision must be made for the tool to collect its **stdin** and pass it to the host RM (via
4 the PMIx server) for forwarding. Two methods of support for forwarding of **stdin** are defined:

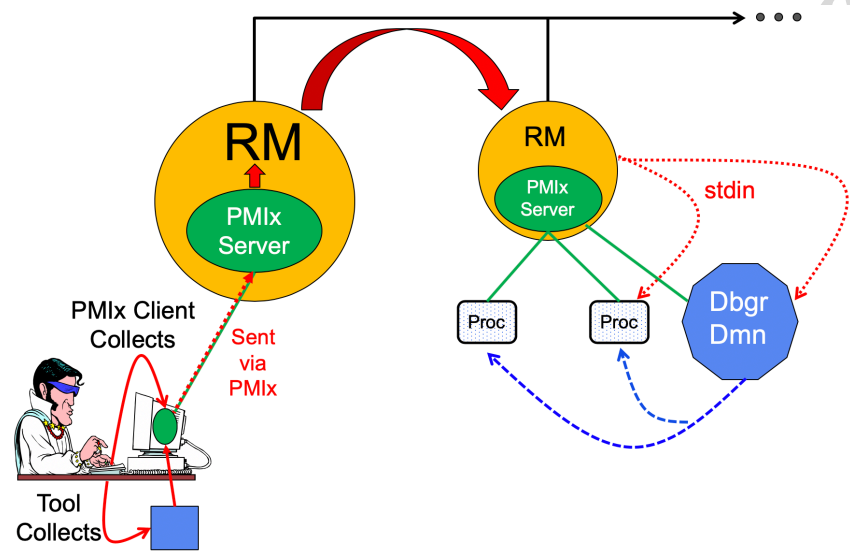


Figure 18.5.: Forwarding stdin

- 5 • internal collection by the PMIx tool library itself. This is requested via the
6 **PMIX_IOF_PUSH_STDIN** attribute in the **PMIx_IOF_push** call. When this mode is
7 selected, the tool library begins collecting all **stdin** data and internally passing it to the local
8 server for distribution to the specified target processes. All collected data is sent to the same
9 targets until **stdin** is closed, or a subsequent call to **PMIx_IOF_push** is made that includes
10 the **PMIX_IOF_COMPLETE** attribute indicating that forwarding of **stdin** is to be terminated.
- 11 • external collection directly by the tool. It is assumed that the tool will provide its own
12 code/mechanism for collecting its **stdin** as the tool developers may choose to insert some
13 filtering and/or editing of the stream prior to forwarding it. In addition, the tool can directly
14 control the targets for the data on a per-call basis – i.e., each call to **PMIx_IOF_push** can
15 specify its own set of target recipients for that particular *blob* of data. Thus, this method provides
16 maximum flexibility, but requires that the tool developer provide their own code to capture
17 **stdin**.

18 Note that it is the responsibility of the RM to forward data to the host where the target process(es)
19 are executing, and for the host daemon on that node to deliver the data to the **stdin** of target
20 process(es). The PMIx server on the remote node is not involved in this process. Systems that do

1 not support forwarding of `stdin` shall return `PMIX_ERR_NOT_SUPPORTED` in response to a
2 forwarding request.

Advice to users

3 Scalable forwarding of `stdin` represents a significant challenge. Most environments will at least
4 handle a *send-to-1* model whereby `stdin` is forwarded to a single identified process, and
5 occasionally an additional *send-to-all* model where `stdin` is forwarded to all processes in the
6 application. Users are advised to check their host environment for available support as the
7 distribution method lies outside the scope of PMIx.

8 `Stdin` buffering by the RM and/or PMIx library can be problematic. If any targeted recipient is
9 slow reading data (or decides never to read data), then the data must be buffered in some
10 intermediate daemon or the PMIx tool library itself. Thus, piping a large amount of data into
11 `stdin` can result in a very large memory footprint in the system management stack or the tool.
12 Best practices, therefore, typically focus on reading of input files by application processes as
13 opposed to forwarding of `stdin`.

14 18.3.3 IO Forwarding Channels

15 *PMIx v3.0*

16 The `pmix_iof_channel_t` structure is a `uint16_t` type that defines a set of bit-mask flags
17 for specifying IO forwarding channels. These can be bitwise OR'd together to reference multiple
18 channels.

```
18 PMIX_FWD_NO_CHANNELS  0x0000  Forward no channels.  
19 PMIX_FWD_STDIN_CHANNEL 0x0001  Forward stdin.  
20 PMIX_FWD_STDOUT_CHANNEL 0x0002  Forward stdout.  
21 PMIX_FWD_STDERR_CHANNEL 0x0004  Forward stderr.  
22 PMIX_FWD_STDDIAG_CHANNEL 0x0008  Forward stddiag, if available.  
23 PMIX_FWD_ALL_CHANNELS 0x00ff  Forward all available channels.
```

24 18.3.4 IO Forwarding constants

```
25 PMIX_ERR_IOF_FAILURE  -172  An IO forwarding operation failed - the affected channel  
26 will be included in the notification.  
27 PMIX_ERR_IOF_COMPLETE -173  IO forwarding of the standard input for this process has  
28 completed - i.e., the stdin file descriptor has closed.
```

1 18.3.5 IO Forwarding attributes

2 The following attributes are used to control IO forwarding behavior at the request of tools. Use of
3 the attributes is optional - any option not provided will revert to some implementation-specific
4 value.

5 **PMIX_IOF_LOCAL_OUTPUT "pmix.iof.local" (bool)**

6 Write output streams to local stdout/err

7 **PMIX_IOF_MERGE_STDERR_STDOUT "pmix.iof.mrg" (bool)**

8 Merge stdout and stderr streams from application procs

9 **PMIX_IOF_CACHE_SIZE "pmix.iof.csize" (uint32_t)**

10 The requested size of the PMIx server cache in bytes for each specified channel. By default,
11 the server is allowed (but not required) to drop all bytes received beyond the max size.

12 **PMIX_IOF_DROP_OLDEST "pmix.iof.old" (bool)**

13 In an overflow situation, the PMIx server is to drop the oldest bytes to make room in the
14 cache.

15 **PMIX_IOF_DROP_NEWEST "pmix.iof.new" (bool)**

16 In an overflow situation, the PMIx server is to drop any new bytes received until room
17 becomes available in the cache (default).

18 **PMIX_IOF_BUFFERING_SIZE "pmix.iof.bsize" (uint32_t)**

19 Requests that IO on the specified channel(s) be aggregated in the PMIx tool library until the
20 specified number of bytes is collected to avoid being called every time a block of IO arrives.
21 The PMIx tool library will execute the callback and reset the collection counter whenever the
22 specified number of bytes becomes available. Any remaining buffered data will be *flushed* to
23 the callback upon a call to deregister the respective channel.

24 **PMIX_IOF_BUFFERING_TIME "pmix.iof.btime" (uint32_t)**

25 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
26 size, this prevents IO from being held indefinitely while waiting for another payload to arrive.

27 *Provisional* **PMIX_IOF_OUTPUT_RAW "pmix.iof.raw" (bool)**

28 Do not buffer output to be written as complete lines - output characters as the stream delivers
29 them

30 **PMIX_IOF_COMPLETE "pmix.iof.cmp" (bool)**

31 Indicates that the specified IO channel has been closed by the source.

32 **PMIX_IOF_TAG_OUTPUT "pmix.iof.tag" (bool)**

33 Requests that output be prefixed with the nspace,rank of the source and a string identifying
34 the channel (*stdout*, *stderr*, etc.).

35 **PMIX_IOF_TIMESTAMP_OUTPUT "pmix.iof.ts" (bool)**

36 Requests that output be marked with the time at which the data was received by the tool -
37 note that this will differ from the time at which the data was collected from the source.

38 *Provisional* **PMIX_IOF_RANK_OUTPUT "pmix.iof.rank" (bool)**

39 Tag output with the rank it came from

40 **PMIX_IOF_XML_OUTPUT "pmix.iof.xml" (bool)**

41 Requests that output be formatted in XML.

42 **PMIX_IOF_PUSH_STDIN "pmix.iof.stdin" (bool)**

1 Requests that the PMIx library collect the `stdin` of the requester and forward it to the
2 processes specified in the `PMIx_IOF_push` call. All collected data is sent to the same
3 targets until `stdin` is closed, or a subsequent call to `PMIx_IOF_push` is made that
4 includes the `PMIx_IOF_COMPLETE` attribute indicating that forwarding of `stdin` is to be
5 terminated.

6 **PMIx_IOF_COPY** "`pmix.iof.cpy`" (bool)

7 Requests that the host environment deliver a copy of the specified output stream(s) to the
8 tool, letting the stream(s) continue to also be delivered to the default location. This allows the
9 tool to tap into the output stream(s) without redirecting it from its current final destination.

10 **PMIx_IOF_REDIRECT** "`pmix.iof.redir`" (bool)

11 Requests that the host environment intercept the specified output stream(s) and deliver it to
12 the requesting tool instead of its current final destination. This might be used, for example,
13 during a debugging procedure to avoid injection of debugger-related output into the
14 application's results file. The original output stream(s) destination is restored upon
15 termination of the tool.

16 *Provisional*

PMIx_IOF_OUTPUT_TO_FILE "`pmix.iof.file`" (char*)

17 Direct application output into files of form "`<filename>.<nospace>.<rank>.stdout`" (for
18 `stdout`) and "`<filename>.<nospace>.<rank>.stderr`" (for `stderr`). If

19 `PMIx_IOF_MERGE_STDERR_STDOUT` was given, then only the `stdout` file will be
20 created and both streams will be written into it.

21 *Provisional*

PMIx_IOF_OUTPUT_TO_DIRECTORY "`pmix.iof.dir`" (char*)

22 Direct application output into files of form "`<directory>/<nospace>/rank.<rank>/stdout`" (for
23 `stdout`) and "`<directory>/<nospace>/rank.<rank>/stderr`" (for `stderr`). If

24 `PMIx_IOF_MERGE_STDERR_STDOUT` was given, then only the `stdout` file will be
25 created and both streams will be written into it.

26 *Provisional*

PMIx_IOF_FILE_PATTERN "`pmix.iof.fpt`" (bool)

27 Specified output file is to be treated as a pattern and not automatically annotated by `nospace`,
28 `rank`, or other parameters. The pattern can use `%n` for the namespace, and `%r` for the rank
29 wherever those quantities are to be placed. The resulting filename will be appended with
30 `".stdout"` for the `stdout` stream and `".stderr"` for the `stderr` stream. If

31 `PMIx_IOF_MERGE_STDERR_STDOUT` was given, then only the `stdout` file will be
32 created and both streams will be written into it.

33 *Provisional*

PMIx_IOF_FILE_ONLY "`pmix.iof.fonly`" (bool)

34 Output only into designated files - do not also output a copy to the console's `stdout/stderr`

35 18.4 Debugger Support

36 Debuggers are a class of tool that merits special consideration due to their particular requirements
37 for access to job-related information and control over process execution. The primary advantage of
38 using PMIx for these purposes lies in the resulting portability of the debugger as it can be used with
39 any system and/or programming model that supports PMIx. In addition to the general tool support
40 described above, debugger support includes:

- 1 • Co-location, co-spawn, and communication wireup of debugger daemons for scalable launch.
2 This includes providing debugger daemons with endpoint connection information across the
3 daemons themselves.
- 4 • Identification of the job that is to be debugged. This includes automatically providing debugger
5 daemons with the job-level information for their target job.

6 Debuggers can also utilize the options in the **PMIx_Spawn** API to exercise a degree of control
7 over spawned jobs for debugging purposes. For example, a debugger can utilize the environmental
8 parameter attributes of Section 12.2.4 to request **LD_PRELOAD** of a memory interceptor library
9 prior to spawning an application process, or interject a custom fork/exec agent to shepherd the
10 application process.

11 A key element of the debugging process is the ability of the debugger to require that processes
12 *pause* at some well-defined point, thereby providing the debugger with an opportunity to attach and
13 control execution. The actual implementation of the *pause* lies outside the scope of PMIx - it
14 typically requires either the launcher or the application itself to implement the necessary
15 operations. However, PMIx does provide several standard attributes by which the debugger can
16 specify the desired attach point:

- 17 • **PMIX_DEBUG_STOP_ON_EXEC** "pmix.dbg.exec" (bool)
18 Included in either the **pmix_info_t** array in a **pmix_app_t** description (if the
19 directive applies only to that application) or in the *job_info* array if it applies to all
20 applications in the given spawn request. Indicates that the application is being spawned
21 under a debugger, and that the local launch agent is to pause the resulting application
22 processes on first instruction for debugger attach. The launcher (RM or IL) is to generate
23 the **PMIX_LAUNCH_COMPLETE** event when all processes are stopped at the exec point.
24 Launchers that cannot support this operation shall return an error from the **PMIx_Spawn**
25 API if this behavior is requested.
- 26 • **PMIX_DEBUG_STOP_IN_INIT** "pmix.dbg.init" (bool)
27 Included in either the **pmix_info_t** array in a **pmix_app_t** description (if the
28 directive applies only to that application) or in the *job_info* array if it applies to all
29 applications in the given spawn request. Indicates that the specified application is being
30 spawned under a debugger. The PMIx client library in each resulting application process
31 shall notify its PMIx server that it is pausing and then pause during **PMIx_Init** of the
32 spawned processes until either released by debugger modification of an appropriate
33 variable or receipt of the **PMIX_DEBUGGER_RELEASE** event. The launcher (RM or IL)
34 is responsible for generating the **PMIX_READY_FOR_DEBUG** event (stipulating a
35 breakpoint of `pmix-init`) when all processes have reached the pause point. PMIx
36 implementations that do not support this operation shall return an error from
37 **PMIx_Init** if this behavior is requested. Launchers that cannot support this operation
38 shall return an error from the **PMIx_Spawn** API if this behavior is requested.
- 39 • **PMIX_DEBUG_STOP_IN_APP** "pmix.dbg.notify" (varies)

1 Direct specified ranks to stop at application-specific point and notify they are
2 ready-to-debug. The attribute's value can be any of three data types:

- 3 – bool - true indicating all ranks
- 4 – `pmix_rank_t` - the rank of one proc, or `PMIX_RANK_WILDCARD` for all
- 5 – a `pmix_data_array_t` if an array of individual processes are specified

6 The resulting application processes are to notify their server (by generating the
7 `PMIX_READY_FOR_DEBUG` event) when they reach some application-determined
8 location - the event shall include the `PMIX_BREAKPOINT` attribute indicating where the
9 application has stopped. The application shall pause at that point until released by
10 debugger modification of an appropriate variable. The launcher (RM or IL) is responsible
11 for generating the `PMIX_READY_FOR_DEBUG` event when all processes have indicated
12 they are at the pause point. Launchers that cannot support this operation shall return an
13 error from the `PMIx_Spawn` API if this behavior is requested.

14 Note that there is no mechanism by which the PMIx library or the launcher can verify that
15 an application will recognize and support the `PMIX_DEBUG_STOP_IN_APP` request.
16 Debuggers utilizing this attachment method must, therefore, be prepared to deal with the
17 case where the application fails to recognize and/or honor the request.

18 If the PMIx implementation and/or the host environment support it, debuggers can utilize the
19 `PMIx_Query_info` API to determine which features are available via the
20 `PMIX_QUERY_ATTRIBUTE_SUPPORT` attribute.

- 21 • `PMIX_DEBUG_STOP_IN_INIT` by checking `PMIX_CLIENT_ATTRIBUTES` for the
22 `PMIx_Init` API.
- 23 • `PMIX_DEBUG_STOP_ON_EXEC` by checking `PMIX_HOST_ATTRIBUTES` for the
24 `PMIx_Spawn` API.

25 The target namespace or process (as given by the debugger in the spawn request) shall be provided
26 to each daemon in its job-level information via the `PMIX_DEBUG_TARGET` attribute. Debugger
27 daemons are responsible for self-determining their specific target process(es), and can then utilize
28 the `PMIx_Query_info` API to obtain information about them (see Fig 18.6) - e.g., to obtain the
29 PIDs of the local processes to which they need to attach. PMIx provides the
30 `pmix_proc_info_t` structure for organizing information about a process' PID, location, and
31 state. Debuggers may request information on a given job at two levels:

- 32 • `PMIX_QUERY_PROC_TABLE` "`pmix.qry.ptable`" (`char*`)
33 Returns a (`pmix_data_array_t`) array of `pmix_proc_info_t`, one entry for each
34 process in the specified namespace, ordered by process job rank. REQUIRED
35 QUALIFIER: `PMIX_NAMESPACE` indicating the namespace whose process table is being
36 queried.
- 37 • `PMIX_QUERY_LOCAL_PROC_TABLE` "`pmix.qry.lptable`" (`char*`)

1 Returns a (`pmix_data_array_t`) array of `pmix_proc_info_t`, one entry for each
2 process in the specified namespace executing on the same node as the requester, ordered
3 by process job rank. REQUIRED QUALIFIER: `PMIX_NAMESPACE` indicating the
4 namespace whose local process table is being queried. OPTIONAL QUALIFIER:
5 `PMIX_HOSTNAME` indicating the host whose local process table is being queried. By
6 default, the query assumes that the host upon which the request was made is to be used.

7 Note that the information provided in the returned proctable represents a snapshot in time. Any
8 process, regardless of role (tool, client, debugger, etc.) can obtain the proctable of a given
9 namespace so long as it has the system-determined authorizations to do so. The list of namespaces
10 available via a given server can be obtained using the `PMIx_Query_info` API with the
11 `PMIX_QUERY_NAMESPACES` key.

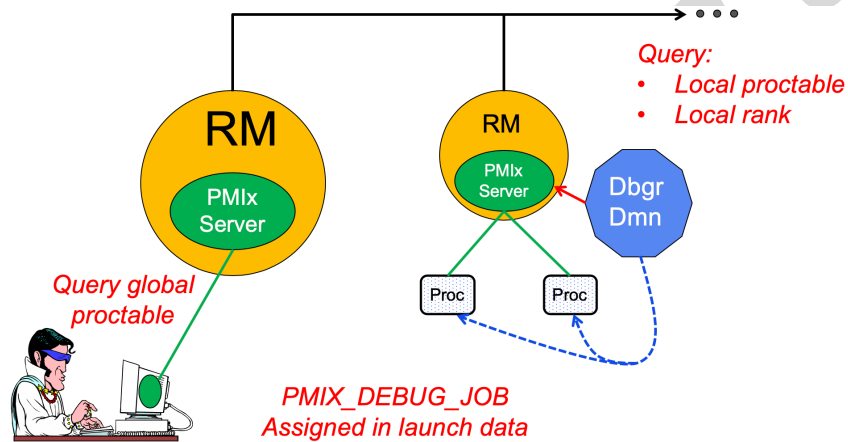


Figure 18.6.: Obtaining proctables

12 Debugger daemons can be started in two ways - either at the same time the application is spawned,
13 or separately at a later time.

14 18.4.1 Co-Location of Debugger Daemons

15 Debugging operations typically require the use of daemons that are located on the same node as the
16 processes they are attempting to debug. The debugger can, of course, specify its own mapping
17 method when issuing its spawn request or utilize its own internal launcher to place the daemons.
18 However, when attaching to a running job, PMIx provides debuggers with a simplified method for
19 requesting that the launcher associated with the job *co-locate* the required daemons. Debuggers can
20 request *co-location* of their daemons by adding the following attributes to the `PMIx_Spawn` used
21 to spawn them:

- 22 • `PMIX_DEBUGGER_DAEMONS` - indicating that the launcher is being asked to spawn debugger
23 daemons.

- 1 • **PMIX_DEBUG_TARGET** - indicating the job or process that is to be debugged. This allows the
2 launcher to identify the processes to be debugged and their location. Note that the debugger job
3 shall be assigned its own namespace (different from that of the job it is being spawned to debug)
4 and each daemon will be assigned a unique rank within that namespace.
- 5 • **PMIX_DEBUG_DAEMONS_PER_PROC** - specifies the number of debugger daemons to be
6 co-located per target process.
- 7 • **PMIX_DEBUG_DAEMONS_PER_NODE** - specifies the number of debugger daemons to be
8 co-located per node where at least one target process is executing.

9 Debugger daemons spawned in this manner shall be provided with the typical PMIx information for
10 their own job plus the target they are to debug via the **PMIX_DEBUG_TARGET** attribute. The
11 debugger daemons spawned on a given node are responsible for self-determining their specific
12 target process(es) - e.g., by referencing their own **PMIX_LOCAL_RANK** in the daemon debugger
13 job versus the corresponding **PMIX_LOCAL_RANK** of the target processes on the node. Note that
14 the debugger will be attaching to the application processes at some arbitrary point in the
15 application's execution unless some method for pausing the application (e.g., by providing a PMIx
16 directive at time of launch, or via a tool using the **PMIx_Job_control** API to direct that the
17 process be paused) has been employed.

Advice to users

18 Note that the tool calling **PMIx_Spawn** to request the launch of the debugger daemons is *not*
19 included in the resulting job - i.e., the debugger daemons do not inherit the namespace of the tool.
20 Thus, collective operations and notifications that target the debugger daemon job will not include
21 the tool unless the namespace/rank of the tool is explicitly included.

22 18.4.2 Co-Spawn of Debugger Daemons

23 In the case where a job is being spawned under the control of a debugger, PMIx provides a shortcut
24 method for spawning the debugger's daemons in parallel with the job. This requires that the
25 debugger be specified as one of the **pmix_app_t** in the same spawn command used to start the
26 job. The debugger application must include at least the **PMIX_DEBUGGER_DAEMONS** attribute
27 identifying itself as a debugger, and may utilize either a mapping option to direct daemon
28 placement, or one of the **PMIX_DEBUG_DAEMONS_PER_PROC** or
29 **PMIX_DEBUG_DAEMONS_PER_NODE** directives.

30 The launcher must not include information regarding the debugger daemons in the job-level info
31 provided to the rest of the **pmix_app_t**s, nor in any calculated rank values (e.g.,
32 **PMIX_NODE_RANK** or **PMIX_LOCAL_RANK**) in those applications. The debugger job is to be
33 assigned its own namespace and each debugger daemon shall receive a unique rank - i.e., the
34 debugger application is to be treated as a completely separate PMIx job that is simply being started
35 in parallel with the user's applications. The launcher is free to implement the launch as a single
36 operation for both the applications and debugger daemons (preferred), or may stage the launches as

1 required. The launcher shall not return from the `PMIx_Spawn` command until all included
2 applications and the debugger daemons have been started.

3 Attributes that apply to both the debugger daemons and the application processes can be specified
4 in the `job_info` array passed into the `PMIx_Spawn` API. Attributes that either (a) apply solely to
5 the debugger daemons or to one of the applications included in the spawn request, or (b) have
6 values that differ from those provided in the `job_info` array, should be specified in the `info` array in
7 the corresponding `pmix_app_t`. Note that PMIx job *pause* attributes (e.g.,
8 `PMIX_DEBUG_STOP_IN_INIT`) do not apply to applications (defined in `pmix_app_t`) where
9 the `PMIX_DEBUGGER_DAEMONS` attribute is set to `true`.

10 Debugger daemons spawned in this manner shall be provided with the typical PMIx information for
11 their own job plus the target they are to debug via the `PMIX_DEBUG_TARGET` attribute. The
12 debugger daemons spawned on a given node are responsible for self-determining their specific
13 target process(es) - e.g., by referencing their own `PMIX_LOCAL_RANK` in the daemon debugger
14 job versus the corresponding `PMIX_LOCAL_RANK` of the target processes on the node.

Advice to users

15 Note that the tool calling `PMIx_Spawn` to request the launch of the debugger daemons is *not*
16 included in the resulting job - i.e., the debugger daemons do not inherit the namespace of the tool.
17 Thus, collective operations and notifications that target the debugger daemon job will not include
18 the tool unless the namespace/rank of the tool is explicitly included.

19 The `PMIx_Spawn` API only supports the return of a single namespace resulting from the spawn
20 request. In the case where the debugger job is co-spawned with the application, the spawn function
21 shall return the namespace of the application and not the debugger job. Tools requiring access to
22 the namespace of the debugger job must query the launcher for the spawned namespaces to find the
23 one belonging to the debugger job.

24 18.4.3 Debugger Agents

25 Individual debuggers may, depending upon implementation, require varying degrees of control over
26 each application process when it is started beyond those available via directives to `PMIx_Spawn`.
27 PMIx offers two mechanisms to help provide a means of meeting these needs.

28 The `PMIX_FORKEXEC_AGENT` attribute allows the debugger to specify an intermediate process
29 (the Fork/Exec Agent (FEA)) for spawning the actual application process (see Fig. 18.7a), thereby
30 interposing the debugger daemon between the application process and the launcher's daemon.
31 Instead of spawning the application process, the launcher will spawn the FEA, which will connect
32 back to the PMIx server as a tool to obtain the spawn description of the application process it is to
33 spawn. The PMIx server in the launcher's daemon shall not register the fork/exec agent as a local
34 client process, nor shall the launcher include the agent in any of the job-level values (e.g.,
35 `PMIX_RANK` within the job or `PMIX_LOCAL_RANK` on the node) provided to the application

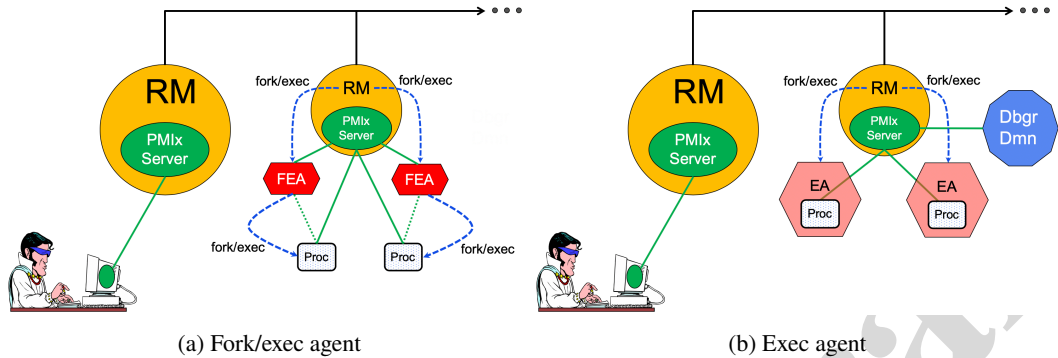


Figure 18.7.: Intermediate agents

1 process. The launcher shall treat the collection of FEAs as a debugger job equivalent to the
 2 co-spawn use-case described in Section 18.4.2.

3 In contrast, the `PMIX_EXEC_AGENT` attribute (Fig. 18.7b) allows the debugger to specify an agent
 4 that will perform some preparatory actions and then exec the eventual application process to replace
 5 itself. In this scenario, the exec agent is provided with the application process' command line as
 6 arguments on its command line (e.g., `"/agent appargv[0] appargv[1]"`) and does not
 7 connect back to the host's PMIx server. It is the responsibility of the exec agent to properly separate
 8 its own command line arguments (if any) from the application description.

9 18.4.4 Tracking the job lifecycle

10 There are a wide range of events a debugger can register to receive, but three are specifically
 11 defined for tracking a job's progress:

- 12 • `PMIX_EVENT_JOB_START` indicates when the first process in the job has been spawned.
- 13 • `PMIX_LAUNCH_COMPLETE` indicates when the last process in the job has been spawned.
- 14 • `PMIX_EVENT_JOB_END` indicates that all processes have terminated.

15 Each event is required to contain at least the namespace of the corresponding job and a
 16 `PMIX_EVENT_TIMESTAMP` indicating the time the event occurred. In addition, the
 17 `PMIX_EVENT_JOB_END` event shall contain the returned status code
 18 (`PMIX_JOB_TERM_STATUS`) for the corresponding job, plus the identity (`PMIX_PROCID`) and
 19 exit status (`PMIX_EXIT_CODE`) of the first failed process, if applicable. Generation of these
 20 events by the launcher can be requested by including the `PMIX_NOTIFY_JOB_EVENTS`
 21 attributes in the spawn request. Note that these events can be logged via the `PMIx_Log` API by
 22 including the `PMIX_LOG_JOB_EVENTS` attribute - this can be done either in conjunction with
 23 generated events, or in place of them.

1 Alternatively, if the debugger or tool solely wants to be alerted to job termination, then including
2 the **PMIX_NOTIFY_COMPLETION** attribute in the spawn request would suffice. This attribute
3 directs the launcher to provide just the **PMIX_EVENT_JOB_END** event. Note that this event can be
4 logged via the **PMIx_Log** API by including the **PMIX_LOG_COMPLETION** attribute - this can be
5 done either in conjunction with the generated event, or in place of it.

Advice to users

6 The PMIx server is required to cache events in order to avoid race conditions - e.g., when a tool is
7 trying to register for the **PMIX_EVENT_JOB_END** event from a very short-lived job. Accordingly,
8 registering for job-related events can result in receiving events relating to jobs other than the one of
9 interest.

10 Users are therefore advised to specify the job whose events are of interest by including the
11 **PMIX_EVENT_AFFECTED_PROC** or **PMIX_EVENT_AFFECTED_PROCS** attribute in the *info*
12 array passed to the **PMIx_Register_event_handler** API.

13 18.4.4.1 Job lifecycle events

14 **PMIX_EVENT_JOB_START -191** The first process in the job has been spawned - includes
15 **PMIX_EVENT_TIMESTAMP** as well as the **PMIX_JOBID** and/or **PMIX_NAMESPACE** of the job.

16 **PMIX_LAUNCH_COMPLETE -174** All processes in the job have been spawned - includes
17 **PMIX_EVENT_TIMESTAMP** as well as the **PMIX_JOBID** and/or **PMIX_NAMESPACE** of the job.

18 **PMIX_EVENT_JOB_END -145** All processes in the job have terminated - includes
19 **PMIX_EVENT_TIMESTAMP** when the last process terminated as well as the **PMIX_JOBID**
20 and/or **PMIX_NAMESPACE** of the job.

21 **PMIX_EVENT_SESSION_START -192** The allocation has been instantiated and is ready
22 for use - includes **PMIX_EVENT_TIMESTAMP** as well as the **PMIX_SESSION_ID** of the
23 allocation. This event is issued after any system-controlled prologue has completed, but
24 before any user-specified actions are taken.

25 **PMIX_EVENT_SESSION_END -193** The allocation has terminated - includes
26 **PMIX_EVENT_TIMESTAMP** as well as the **PMIX_SESSION_ID** of the allocation. This
27 event is issued after any user-specified actions have completed, but before any
28 system-controlled epilogue is performed.

29 The following events relate to processes within a job:

30 **PMIX_EVENT_PROC_TERMINATED -201** The specified process(es) terminated - normal
31 or abnormal termination will be indicated by the **PMIX_PROC_TERM_STATUS** in the *info*
32 array of the notification. Note that a request for individual process events can generate a
33 significant event volume from large-scale jobs.

34 **PMIX_ERR_PROC_TERM_WO_SYNC -200** Process terminated without calling
35 **PMIx_Finalize**, or was a member of an assemblage formed via **PMIx_Connect** and
36 terminated or called **PMIx_Finalize** without first calling **PMIx_Disconnect** (or its
37 non-blocking form) from that assemblage.

The following constants may be included via the `PMIX_JOB_TERM_STATUS` attributed in the `info` array in the `PMIX_EVENT_JOB_END` event notification to provide more detailed information regarding the reason for job abnormal termination:

`PMIX_ERR_JOB_CANCELED` -180 The job was canceled by the host environment.
`PMIX_ERR_JOB_ABORTED` -182 One or more processes in the job called abort, causing the job to be terminated.
`PMIX_ERR_JOB_KILLED_BY_CMD` -183 The job was killed by user command.
`PMIX_ERR_JOB_ABORTED_BY_SIG` -184 The job was aborted due to receipt of an error signal (e.g., SIGKILL).
`PMIX_ERR_JOB_TERM_WO_SYNC` -185 The job was terminated due to at least one process terminating without calling `PMIx_Finalize`, or was a member of an assemblage formed via `PMIx_Connect` and terminated or called `PMIx_Finalize` without first calling `PMIx_Disconnect` (or its non-blocking form) from that assemblage.
`PMIX_ERR_JOB_SENSOR_BOUND_EXCEEDED` -186 The job was terminated due to one or more processes exceeding a specified sensor limit.
`PMIX_ERR_JOB_NON_ZERO_TERM` -187 The job was terminated due to one or more processes exiting with a non-zero status.
`PMIX_ERR_JOB_ABORTED_BY_SYS_EVENT` -189 The job was aborted due to receipt of a system event.

18.4.4.2 Job lifecycle attributes

`PMIX_JOB_TERM_STATUS` "pmix.job.term.status" (`pmix_status_t`)

Status returned by job upon its termination. The status will be communicated as part of a PMIx event payload provided by the host environment upon termination of a job. Note that generation of the `PMIX_EVENT_JOB_END` event is optional and host environments may choose to provide it only upon request.

`PMIX_PROC_STATE_STATUS` "pmix.proc.state" (`pmix_proc_state_t`)

State of the specified process as of the last report - may not be the actual current state based on update rate.

`PMIX_PROC_TERM_STATUS` "pmix.proc.term.status" (`pmix_status_t`)

Status returned by a process upon its termination. The status will be communicated as part of a PMIx event payload provided by the host environment upon termination of a process. Note that generation of the `PMIX_EVENT_PROC_TERMINATED` event is optional and host environments may choose to provide it only upon request.

18.4.5 Debugger-related constants

The following constants are used in events used to coordinate applications and the debuggers attaching to them.

`PMIX_READY_FOR_DEBUG` -58 Event indicating a job (or specified set of processes) is ready for debug - includes identification of the target processes as well as the

`PMIX_BREAKPOINT` indicating where the target is waiting

`PMIX_DEBUGGER_RELEASE` -3 Release a tool that is paused during `PMIx_tool_init`.

1 18.4.6 Debugger attributes

2 Attributes used to assist debuggers - these are values that can either be passed to the `PMIx_Spawn`
3 APIs or accessed by a debugger itself using the `PMIx_Get` API with the
4 `PMIX_RANK_WILDCARD` rank.

5 **PMIX_DEBUG_STOP_ON_EXEC** "pmix.dbg.exec" (bool)

6 Included in either the `pmix_info_t` array in a `pmix_app_t` description (if the directive
7 applies only to that application) or in the `job_info` array if it applies to all applications in the
8 given spawn request. Indicates that the application is being spawned under a debugger, and
9 that the local launch agent is to pause the resulting application processes on first instruction
10 for debugger attach. The launcher (RM or IL) is to generate the
11 `PMIX_LAUNCH_COMPLETE` event when all processes are stopped at the exec point.

12 **PMIX_DEBUG_STOP_IN_INIT** "pmix.dbg.init" (bool)

13 Included in either the `pmix_info_t` array in a `pmix_app_t` description (if the directive
14 applies only to that application) or in the `job_info` array if it applies to all applications in the
15 given spawn request. Indicates that the specified application is being spawned under a
16 debugger. The PMIx client library in each resulting application process shall notify its PMIx
17 server that it is pausing and then pause during `PMIx_Init` of the spawned processes until
18 either released by debugger modification of an appropriate variable or receipt of the
19 `PMIX_DEBUGGER_RELEASE` event. The launcher (RM or IL) is responsible for generating
20 the `PMIX_READY_FOR_DEBUG` event (stipulating a breakpoint of `pmix-init`) when all
21 processes have reached the pause point.

22 **PMIX_DEBUG_STOP_IN_APP** "pmix.dbg.notify" (varies)

23 Direct specified ranks to stop at application-specific point and notify they are
24 ready-to-debug. The attribute's value can be any of three data types:

- 25 • bool - true indicating all ranks
- 26 • `pmix_rank_t` - the rank of one proc, or `PMIX_RANK_WILDCARD` for all
- 27 • a `pmix_data_array_t` if an array of individual processes are specified

28 The resulting application processes are to notify their server (by generating the
29 `PMIX_READY_FOR_DEBUG` event) when they reach some application-determined location
30 - the event shall include the `PMIX_BREAKPOINT` attribute indicating where the application
31 has stopped. The application shall pause at that point until released by debugger
32 modification of an appropriate variable. The launcher (RM or IL) is responsible for
33 generating the `PMIX_READY_FOR_DEBUG` event when all processes have indicated they
34 are at the pause point.

35 **PMIX_BREAKPOINT** "pmix.brkpt" (char*)

36 String ID of the breakpoint where the process(es) is(are) waiting.

37 **PMIX_DEBUG_TARGET** "pmix.dbg.tgt" (`pmix_proc_t*`)

38 Identifier of process(es) to be debugged - a rank of `PMIX_RANK_WILDCARD` indicates that
39 all processes in the specified namespace are to be included.

40 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)

41 Included in the `pmix_info_t` array of a `pmix_app_t`, this attribute declares that the
42 application consists of debugger daemons and shall be governed accordingly. If used as the
43 sole `pmix_app_t` in a `PMIx_Spawn` request, then the `PMIX_DEBUG_TARGET` attribute

1 must also be provided (in either the *job_info* or in the *info* array of the `pmix_app_t`) to
2 identify the namespace to be debugged so that the launcher can determine where to place the
3 spawned daemons. If neither `PMIX_DEBUG_DAEMONS_PER_PROC` nor
4 `PMIX_DEBUG_DAEMONS_PER_NODE` is specified, then the launcher shall default to a
5 placement policy of one daemon per process in the target job.

6 **PMIX_COSPAWN_APP** "`pmix.cospawn`" (bool)

7 Designated application is to be spawned as a disconnected job - i.e., the launcher shall not
8 include the application in any of the job-level values (e.g., `PMIX_RANK` within the job)
9 provided to any other application process generated by the same spawn request. Typically
10 used to cospawn debugger daemons alongside an application.

11 **PMIX_DEBUG_DAEMONS_PER_PROC** "`pmix.dbg.dpproc`" (uint16_t)

12 Number of debugger daemons to be spawned per application process. The launcher is to pass
13 the identifier of the namespace to be debugged by including the `PMIX_DEBUG_TARGET`
14 attribute in the daemon's job-level information. The debugger daemons spawned on a given
15 node are responsible for self-determining their specific target process(es) - e.g., by
16 referencing their own `PMIX_LOCAL_RANK` in the daemon debugger job versus the
17 corresponding `PMIX_LOCAL_RANK` of the target processes on the node.

18 **PMIX_DEBUG_DAEMONS_PER_NODE** "`pmix.dbg.dpnd`" (uint16_t)

19 Number of debugger daemons to be spawned on each node where the target job is executing.
20 The launcher is to pass the identifier of the namespace to be debugged by including the
21 `PMIX_DEBUG_TARGET` attribute in the daemon's job-level information. The debugger
22 daemons spawned on a given node are responsible for self-determining their specific target
23 process(es) - e.g., by referencing their own `PMIX_LOCAL_RANK` in the daemon debugger
24 job versus the corresponding `PMIX_LOCAL_RANK` of the target processes on the node.

25 **PMIX_QUERY_PROC_TABLE** "`pmix.qry.phtable`" (char*)

26 Returns a (`pmix_data_array_t`) array of `pmix_proc_info_t`, one entry for each
27 process in the specified namespace, ordered by process job rank. REQUIRED QUALIFIER:
28 `PMIX_NAMESPACE` indicating the namespace whose process table is being queried.

29 **PMIX_QUERY_LOCAL_PROC_TABLE** "`pmix.qry.lhtable`" (char*)

30 Returns a (`pmix_data_array_t`) array of `pmix_proc_info_t`, one entry for each
31 process in the specified namespace executing on the same node as the requester, ordered by
32 process job rank. REQUIRED QUALIFIER: `PMIX_NAMESPACE` indicating the namespace
33 whose local process table is being queried. OPTIONAL QUALIFIER: `PMIX_HOSTNAME`
34 indicating the host whose local process table is being queried. By default, the query assumes
35 that the host upon which the request was made is to be used.

36 18.5 Tool-Specific APIs

37 PMIx-based tools automatically have access to all PMIx client functions. Tools designated as a
38 *launcher* or a *server* will also have access to all PMIx server functions. There are, however, an
39 additional set of functions (described in this section) that are specific to a PMIx tool. Access to
40 those functions require use of the tool initialization routine.

1 18.5.1 PMIx_tool_init

2 Summary

3 Initialize the PMIx library for operating as a tool, optionally connecting to a specified PMIx server.

4 Format

PMIx v2.0

```
pmix_status_t  
PMIx_tool_init(pmix_proc_t *proc,  
               pmix_info_t info[], size_t ninfo);
```

8 INOUT proc

9 `pmix_proc_t` structure (handle)

10 IN info

11 Array of `pmix_info_t` structures (array of handles)

12 IN ninfo

13 Number of elements in the *info* array (`size_t`)

14 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

15 The following attributes are required to be supported by all PMIx libraries:

16 **PMIX_TOOL_NAMESPACE** "pmix.tool.namespace" (`char*`)

17 Name of the namespace to use for this tool.

18 **PMIX_TOOL_RANK** "pmix.tool.rank" (`uint32_t`)

19 Rank of this tool.

20 **PMIX_TOOL_DO_NOT_CONNECT** "pmix.tool.nocon" (`bool`)

21 The tool wants to use internal PMIx support, but does not want to connect to a PMIx server.

22 **PMIX_TOOL_ATTACHMENT_FILE** "pmix.tool.attach" (`char*`)

23 Pathname of file containing connection information to be used for attaching to a specific
24 server.

25 **PMIX_SERVER_URI** "pmix.srvr.uri" (`char*`)

26 URI of the PMIx server to be contacted.

27 **PMIX_TCP_URI** "pmix.tcp.uri" (`char*`)

28 The URI of the PMIx server to connect to, or a file name containing it in the form of
29 `file:<name of file containing it>`.

30 **PMIX_SERVER_PIDINFO** "pmix.srvr.pidinfo" (`pid_t`)

31 PID of the target PMIx server for a tool.

32 **PMIX_SERVER_NAMESPACE** "pmix.srv.namespace" (`char*`)

33 Name of the namespace to use for this PMIx server.

1 **PMIX_CONNECT_TO_SYSTEM** "pmix.cnct.sys" (bool)

2 The requester requires that a connection be made only to a local, system-level PMIx server.

3 **PMIX_CONNECT_SYSTEM_FIRST** "pmix.cnct.sys.first" (bool)

4 Preferentially, look for a system-level PMIx server first.

▲-----▲
▼-----▼ **Optional Attributes** -----▼

5 The following attributes are optional for implementers of PMIx libraries:

6 **PMIX_CONNECT_RETRY_DELAY** "pmix.tool.retry" (uint32_t)

7 Time in seconds between connection attempts to a PMIx server - the default value is
8 implementation specific.

9 **PMIX_CONNECT_MAX_RETRIES** "pmix.tool.mretries" (uint32_t)

10 Maximum number of times to try to connect to PMIx server - the default value is
11 implementation specific.

12 **PMIX_SOCKET_MODE** "pmix.sockmode" (uint32_t)

13 POSIX *mode_t* (9 bits valid). If the library supports socket connections, this attribute may
14 be supported for setting the socket mode.

15 **PMIX_TCP_REPORT_URI** "pmix.tcp.repuri" (char*)

16 If provided, directs that the TCP URI be reported and indicates the desired method of
17 reporting: '-' for stdout, '+' for stderr, or filename. If the library supports TCP socket
18 connections, this attribute may be supported for reporting the URI.

19 **PMIX_TCP_IF_INCLUDE** "pmix.tcp.ifinclude" (char*)

20 Comma-delimited list of devices and/or CIDR notation to include when establishing the
21 TCP connection. If the library supports TCP socket connections, this attribute may be
22 supported for specifying the interfaces to be used.

23 **PMIX_TCP_IF_EXCLUDE** "pmix.tcp.ifexclude" (char*)

24 Comma-delimited list of devices and/or CIDR notation to exclude when establishing the
25 TCP connection. If the library supports TCP socket connections, this attribute may be
26 supported for specifying the interfaces that are *not* to be used.

27 **PMIX_TCP_IPV4_PORT** "pmix.tcp.ipv4" (int)

28 The IPv4 port to be used.. If the library supports IPV4 connections, this attribute may be
29 supported for specifying the port to be used.

30 **PMIX_TCP_IPV6_PORT** "pmix.tcp.ipv6" (int)

31 The IPv6 port to be used. If the library supports IPV6 connections, this attribute may be
32 supported for specifying the port to be used.

33 **PMIX_TCP_DISABLE_IPV4** "pmix.tcp.disipv4" (bool)

34 Set to **true** to disable IPv4 family of addresses. If the library supports IPV4 connections,
35 this attribute may be supported for disabling it.

1 **PMIX_TCP_DISABLE_IPV6** "pmix.tcp.disipv6" (bool)
 2 Set to **true** to disable IPv6 family of addresses. If the library supports IPV6 connections,
 3 this attribute may be supported for disabling it.

4 **PMIX_EXTERNAL_PROGRESS** "pmix.evect" (bool)
 5 The host shall progress the PMIx library via calls to **PMIx_Progress**

6 **PMIX_EVENT_BASE** "pmix.evbase" (void*)
 7 Pointer to an **event_base** to use in place of the internal progress thread. All PMIx library
 8 events are to be assigned to the provided event base. The event base *must* be compatible with
 9 the event library used by the PMIx implementation - e.g., either both the host and PMIx
 10 library must use libevent, or both must use libev. Cross-matches are unlikely to work and
 11 should be avoided - it is the responsibility of the host to ensure that the PMIx
 12 implementation supports (and was built with) the appropriate event library.

13 **PMIX_IOF_LOCAL_OUTPUT** "pmix.iof.local" (bool)
 14 Write output streams to local stdout/err

15 Description

16 Initialize the PMIx tool, returning the process identifier assigned to this tool in the provided
 17 **pmix_proc_t** struct. The *info* array is used to pass user requests pertaining to the initialization
 18 and subsequent operations. Passing a **NULL** value for the array pointer is supported if no directives
 19 are desired.

20 If called with the **PMIX_TOOL_DO_NOT_CONNECT** attribute, the PMIx tool library will fully
 21 initialize but not attempt to connect to a PMIx server. The tool can connect to a server at a later
 22 point in time, if desired, by calling the **PMIx_tool_attach_to_server** function. If provided,
 23 the *proc* structure will be set to a zero-length namespace and a rank of **PMIX_RANK_UNDEF** unless
 24 the **PMIX_TOOL_NAMESPACE** and **PMIX_TOOL_RANK** attributes are included in the *info* array.

25 In all other cases, the PMIx tool library will automatically attempt to connect to a PMIx server
 26 according to the precedence chain described in Section 18.1. If successful, the function will return
 27 **PMIX_SUCCESS** and will fill the process structure (if provided) with the assigned namespace and
 28 rank of the tool. The server to which the tool connects will be designated its *primary* server. Note
 29 that each connection attempt in the above precedence chain will retry (with delay between each
 30 retry) a number of times according to the values of the corresponding attributes.

31 Note that the PMIx tool library is referenced counted, and so multiple calls to **PMIx_tool_init**
 32 are allowed. If the tool is not connected to any server when this API is called, then the tool will
 33 attempt to connect to a server unless the **PMIX_TOOL_DO_NOT_CONNECT** is included in the call
 34 to API.

35 18.5.2 PMIx_tool_finalize

36 Summary

37 Finalize the PMIx tool library.

1 *PMIx v2.0*

Format

C

2 `pmix_status_t`

3 `PMIx_tool_finalize(void);`

C

4 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Description

6 Finalize the PMIx tool library, closing all existing connections to servers. An error code will be
7 returned if, for some reason, a connection cannot be cleanly terminated — in such cases, the
8 connection is dropped. Upon detecting loss of the connection, the PMIx server shall cleanup all
9 associated records of the tool.

10 18.5.3 `PMIx_tool_disconnect`

11 Summary

12 Disconnect the PMIx tool from the specified server connection while leaving the tool library
13 initialized.

14 *PMIx v4.0* Format

C

15 `pmix_status_t`

16 `PMIx_tool_disconnect(const pmix_proc_t *server);`

C

17 **IN** `server`

18 `pmix_proc_t` structure (handle)

19 Returns `PMIX_SUCCESS` or a negative value indicating the error.

20 Description

21 Close the current connection to the specified server, if one has been made, while leaving the PMIx
22 library initialized. An error code will be returned if, for some reason, the connection cannot be
23 cleanly terminated - in this case, the connection is dropped. In either case, the library will remain
24 initialized. Upon detecting loss of the connection, the PMIx server shall cleanup all associated
25 records of the tool.

26 Note that if the server being disconnected is the current *primary* server, then all operations
27 requiring support from a server will return the `PMIX_ERR_UNREACH` error until the tool either
28 designates an existing connection to be the *primary* server or, if no other connections exist, the tool
29 establishes a connection to a PMIx server.

1 18.5.4 PMix_tool_attach_to_server

2 Summary

3 Establish a connection to a PMIx server.

4 Format

```
5 pmix_status_t  
6 PMIX_tool_attach_to_server(pmix_proc_t *proc,  
7                             pmix_proc_t *server,  
8                             pmix_info_t info[],  
9                             size_t ninfo);
```

10 INOUT proc

11 Pointer to `pmix_proc_t` structure (handle)

12 INOUT server

13 Pointer to `pmix_proc_t` structure (handle)

14 IN info

15 Array of `pmix_info_t` structures (array of handles)

16 IN ninfo

17 Number of elements in the *info* array (`size_t`)

18 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

19 The following attributes are required to be supported by all PMIx libraries:

20 **PMIX_TOOL_ATTACHMENT_FILE** "pmix.tool.attach" (`char*`)

21 Pathname of file containing connection information to be used for attaching to a specific
22 server.

23 **PMIX_SERVER_URI** "pmix.srvr.uri" (`char*`)

24 URI of the PMIx server to be contacted.

25 **PMIX_TCP_URI** "pmix.tcp.uri" (`char*`)

26 The URI of the PMIx server to connect to, or a file name containing it in the form of
27 `file:<name of file containing it>`.

28 **PMIX_SERVER_PIDINFO** "pmix.srvr.pidinfo" (`pid_t`)

29 PID of the target PMIx server for a tool.

30 **PMIX_SERVER_NAMESPACE** "pmix.srv.namespace" (`char*`)

31 Name of the namespace to use for this PMIx server.

32 **PMIX_CONNECT_TO_SYSTEM** "pmix.cnct.sys" (`bool`)

33 The requester requires that a connection be made only to a local, system-level PMIx server.

34 **PMIX_CONNECT_SYSTEM_FIRST** "pmix.cnct.sys.first" (`bool`)

1 Preferentially, look for a system-level PMIx server first.

2 **PMIX_PRIMARY_SERVER** "pmix.pri.srvr" (bool)

3 The server to which the tool is connecting shall be designated the *primary* server once
4 connection has been accomplished.

5 **Description**

6 Establish a connection to a server. This function can be called at any time by a PMIx tool to create a
7 new connection to a server. If a specific server is given and the tool is already attached to it, then
8 the API shall return **PMIX_SUCCESS** without taking any further action. In all other cases, the tool
9 will attempt to discover a server using the method described in Section 18.1, ignoring all candidates
10 to which it is already connected. The **PMIX_ERR_UNREACH** error shall be returned if no new
11 connection is made.

12 The process identifier assigned to this tool is returned in the provided *proc* structure. Passing a
13 value of **NULL** for the *proc* parameter is allowed if the user wishes solely to connect to a PMIx
14 server and does not require return of the identifier at that time.

15 The process identifier of the server to which the tool attached is returned in the *server* structure.
16 Passing a value of **NULL** for the *proc* parameter is allowed if the user wishes solely to connect to a
17 PMIx server and does not require return of the identifier at that time.

18 Note that the **PMIX_PRIMARY_SERVER** attribute must be included in the *info* array if the server
19 being connected to is to become the primary server, or a call to **PMIx_tool_set_server** must
20 be provided immediately after the call to this function.

21 **Advice to PMIx library implementers**

22 When a tool connects to a server that is under a different namespace manager (e.g., host RM) from
23 the prior server, the namespace in the identifier of the tool must remain unique in the new universe.
24 If the namespace of the tool fails to meet this criteria in the new universe, then the new namespace
manager is required to return an error and the connection attempt must fail.

25 **Advice to users**

26 Some PMIx implementations may not support connecting to a server that is not under the same
namespace manager (e.g., host RM) as the server to which the tool is currently connected.

27 **18.5.5 PMIx_tool_get_servers**

28 **Summary**

29 Get an array containing the **pmix_proc_t** process identifiers of all servers to which the tool is
30 currently connected.

1 **Format** C

```
2 pmix_status_t
3 PMIx_tool_get_servers(pmix_proc_t *servers[], size_t *nservers);
C
```

4 **OUT servers**
5 Address where the pointer to an array of `pmix_proc_t` structures shall be returned (handle)

6 **INOUT nservers**
7 Address where the number of elements in *servers* shall be returned (handle)

8 Returns `PMIX_SUCCESS` or a negative value indicating the error.

9 **Description**
10 Return an array containing the `pmix_proc_t` process identifiers of all servers to which the tool is
11 currently connected. The process identifier of the current primary server shall be the first entry in
12 the array, with the remaining entries in order of attachment from earliest to most recent.

13 18.5.6 PMIx_tool_set_server

14 **Summary**
15 Designate a server as the tool's *primary* server.

16 **Format** C

PMIx v4.0

```
17 pmix_status_t
18 PMIx_tool_set_server(const pmix_proc_t *server
19 info[], size_t ninfo);
C
```

`pmix_info_t`

20 **IN server**
21 `pmix_proc_t` structure (handle)

22 **IN info**
23 Array of `pmix_info_t` structures (array of handles)

24 **IN ninfo**
25 Number of elements in the *info* array (`size_t`)

26 Returns `PMIX_SUCCESS` or a negative value indicating the error.

Required Attributes

The following attributes are required to be supported by all PMIx libraries:

PMIX_WAIT_FOR_CONNECTION "pmix.wait.conn" (bool)

Wait until the specified process has connected to the requesting tool or server, or the operation times out (if the **PMIX_TIMEOUT** directive is included in the request).

PMIX_TIMEOUT "pmix.timeout" (int)

Time in seconds before the specified operation should time out (zero indicating infinite) and return the **PMIX_ERR_TIMEOUT** error. Care should be taken to avoid race conditions caused by multiple layers (client, server, and host) simultaneously timing the operation.

Description

Designate the specified server to be the tool's *primary* server for all subsequent API calls.

18.5.7 PMIx_IOF_pull

Summary

Register to receive output forwarded from a set of remote processes.

Format

PMIx v3.0

C

```
pmix_status_t
PMIx_IOF_pull(const pmix_proc_t procs[], size_t nprocs,
              const pmix_info_t directives[], size_t ndirs,
              pmix_iof_channel_t channel,
              pmix_iof_cbfunc_t cbfunc,
              pmix_hdlr_reg_cbfunc_t regcbfunc,
              void *regcbdata);
```

C

IN **procs**

Array of proc structures identifying desired source processes (array of handles)

IN **nprocs**

Number of elements in the *procs* array (integer)

IN **directives**

Array of **pmix_info_t** structures (array of handles)

IN **ndirs**

Number of elements in the *directives* array (integer)

IN **channel**

Bitmask of IO channels included in the request (**pmix_iof_channel_t**)

IN **cbfunc**

Callback function for delivering relevant output (**pmix_iof_cbfunc_t** function reference)

1 **IN regcbfunc**
2 Function to be called when registration is completed (`pmix_hdlr_reg_cbfunc_t`
3 function reference)

4 **IN regcbdata**
5 Data to be passed to the `regcbfunc` callback function (memory reference)

6 Returns `PMIX_SUCCESS` or a negative value indicating the error. In the event the function returns
7 an error, the `regcbfunc` will *not* be called.

▼----- Required Attributes -----▼

8 The following attributes are required for PMIx libraries that support IO forwarding:

9 **PMIX_IOF_CACHE_SIZE** "`pmix.iof.csize`" (`uint32_t`)
10 The requested size of the PMIx server cache in bytes for each specified channel. By default,
11 the server is allowed (but not required) to drop all bytes received beyond the max size.

12 **PMIX_IOF_DROP_OLDEST** "`pmix.iof.old`" (`bool`)
13 In an overflow situation, the PMIx server is to drop the oldest bytes to make room in the
14 cache.

15 **PMIX_IOF_DROP_NEWEST** "`pmix.iof.new`" (`bool`)
16 In an overflow situation, the PMIx server is to drop any new bytes received until room
17 becomes available in the cache (default).

▲-----

▼----- Optional Attributes -----▼

18 The following attributes are optional for PMIx libraries that support IO forwarding:

19 **PMIX_IOF_BUFFERING_SIZE** "`pmix.iof.bsize`" (`uint32_t`)
20 Requests that IO on the specified channel(s) be aggregated in the PMIx tool library until the
21 specified number of bytes is collected to avoid being called every time a block of IO arrives.
22 The PMIx tool library will execute the callback and reset the collection counter whenever the
23 specified number of bytes becomes available. Any remaining buffered data will be *flushed* to
24 the callback upon a call to deregister the respective channel.

25 **PMIX_IOF_BUFFERING_TIME** "`pmix.iof.btime`" (`uint32_t`)
26 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
27 size, this prevents IO from being held indefinitely while waiting for another payload to
28 arrive.

29 **PMIX_IOF_TAG_OUTPUT** "`pmix.iof.tag`" (`bool`)
30 Requests that output be prefixed with the namespace, rank of the source and a string identifying
31 the channel (`stdout`, `stderr`, etc.).

32 **PMIX_IOF_TIMESTAMP_OUTPUT** "`pmix.iof.ts`" (`bool`)
33 Requests that output be marked with the time at which the data was received by the tool -
34 note that this will differ from the time at which the data was collected from the source.

1 **PMIX_IOF_XML_OUTPUT** "pmix.iof.xml" (bool)

2 Requests that output be formatted in XML.

3 **Description**

4 Register to receive output forwarded from a set of remote processes.

▼ Advice to users ▼

5 Providing a **NULL** function pointer for the *cbfunc* parameter will cause output for the indicated
6 channels to be written to their corresponding **stdout/stderr** file descriptors. Use of
7 **PMIX_RANK_WILDCARD** to specify all processes in a given namespace is supported but should be
8 used carefully due to bandwidth and memory footprint considerations.

9 **18.5.8 PMIx_IOF_deregister**

10 **Summary**

11 Deregister from output forwarded from a set of remote processes.

12 **Format**

PMIx v3.0

```
13 pmix_status_t  
14 PMIx_IOF_deregister(size_t iofhdlr,  
15                     const pmix_info_t directives[], size_t ndirs,  
16                     pmix_op_cbfunc_t cbfunc, void *cbdata);
```

17 **IN iofhdlr**

18 Registration number returned from the **pmix_hdlr_reg_cbfunc_t** callback from the
19 call to **PMIx_IOF_pull** (**size_t**)

20 **IN directives**

21 Array of **pmix_info_t** structures (array of handles)

22 **IN ndirs**

23 Number of elements in the *directives* array (integer)

24 **IN cbfunc**

25 Callback function to be called when deregistration has been completed. (function reference)

26 **IN cbdata**

27 Data to be passed to the *cbfunc* callback function (memory reference)

28 A successful return indicates that the request is being processed and the result will be returned in
29 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
30 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

31 Returns **PMIX_SUCCESS** or one of the following error codes when the condition described occurs:

1 **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed
2 successfully - the *cbfunc* will *not* be called.

3 If none of the above return codes are appropriate, then an implementation must return either a
4 general PMIx error code or an implementation defined error code as described in Section 3.1.1.

5 **Description**

6 Deregister from output forwarded from a set of remote processes.

▼ Advice to PMIx library implementers ▼

7 Any currently buffered IO should be flushed upon receipt of a deregistration request. All received
8 IO after receipt of the request shall be discarded.



9 **18.5.9 PMIx_IOF_push**

10 **Summary**

11 Push data collected locally (typically from **stdin** or a file) to **stdin** of the target recipients.

12 **Format**

PMIx v3.0

C

```
13 pmix_status_t  
14 PMIx_IOF_push(const pmix_proc_t targets[], size_t ntargets,  
15               pmix_byte_object_t *bo,  
16               const pmix_info_t directives[], size_t ndirs,  
17               pmix_op_cbfunc_t cbfunc, void *cbdata);
```

C

- 18 **IN targets**
19 Array of proc structures identifying desired target processes (array of handles)
- 20 **IN ntargets**
21 Number of elements in the *targets* array (integer)
- 22 **IN bo**
23 Pointer to **pmix_byte_object_t** containing the payload to be delivered (handle)
- 24 **IN directives**
25 Array of **pmix_info_t** structures (array of handles)
- 26 **IN ndirs**
27 Number of elements in the *directives* array (integer)
- 28 **IN directives**
29 Array of **pmix_info_t** structures (array of handles)
- 30 **IN cbfunc**
31 Callback function to be called when operation has been completed. (**pmix_op_cbfunc_t**
32 function reference)

1 **IN cldata**

2 Data to be passed to the *cbfunc* callback function (memory reference)

3 A successful return indicates that the request is being processed and the result will be returned in
4 the provided *cbfunc*. Note that the library must not invoke the callback function prior to returning
5 from the API. The callback function, *cbfunc*, is only called when **PMIX_SUCCESS** is returned.

6 Returns PMIX_SUCCESS or one of the following error codes when the condition described occurs:

7 **PMIX_OPERATION_SUCCEEDED**, indicating that the request was immediately processed
8 successfully - the *cbfunc* will *not* be called.

9 If none of the above return codes are appropriate, then an implementation must return either a
10 general PMIX error code or an implementation defined error code as described in Section 3.1.1.

▼----- Required Attributes -----▼

11 The following attributes are required for PMIX libraries that support IO forwarding:

12 **PMIX_IOF_CACHE_SIZE** "pmix.iof.csize" (uint32_t)

13 The requested size of the PMIX server cache in bytes for each specified channel. By default,
14 the server is allowed (but not required) to drop all bytes received beyond the max size.

15 **PMIX_IOF_DROP_OLDEST** "pmix.iof.old" (bool)

16 In an overflow situation, the PMIX server is to drop the oldest bytes to make room in the
17 cache.

18 **PMIX_IOF_DROP_NEWEST** "pmix.iof.new" (bool)

19 In an overflow situation, the PMIX server is to drop any new bytes received until room
20 becomes available in the cache (default).

▲----- Optional Attributes -----▼

21 The following attributes are optional for PMIX libraries that support IO forwarding:

22 **PMIX_IOF_BUFFERING_SIZE** "pmix.iof.bsize" (uint32_t)

23 Requests that IO on the specified channel(s) be aggregated in the PMIX tool library until the
24 specified number of bytes is collected to avoid being called every time a block of IO arrives.
25 The PMIX tool library will execute the callback and reset the collection counter whenever the
26 specified number of bytes becomes available. Any remaining buffered data will be *flushed* to
27 the callback upon a call to deregister the respective channel.

28 **PMIX_IOF_BUFFERING_TIME** "pmix.iof.btime" (uint32_t)

29 Max time in seconds to buffer IO before delivering it. Used in conjunction with buffering
30 size, this prevents IO from being held indefinitely while waiting for another payload to
31 arrive.

32 **PMIX_IOF_PUSH_STDIN** "pmix.iof.stdin" (bool)

1 Requests that the PMIx library collect the **stdin** of the requester and forward it to the
2 processes specified in the **PMIx_IOF_push** call. All collected data is sent to the same
3 targets until **stdin** is closed, or a subsequent call to **PMIx_IOF_push** is made that
4 includes the **PMIX_IOF_COMPLETE** attribute indicating that forwarding of **stdin** is to be
5 terminated.

6 **Description**

7 Called either to:

- 8 • push data collected by the caller themselves (typically from **stdin** or a file) to **stdin** of the
9 target recipients;
- 10 • request that the PMIx library automatically collect and push the **stdin** of the caller to the target
11 recipients; or
- 12 • indicate that automatic collection and transmittal of **stdin** is to stop

▼ **Advice to users** ▼

13 Execution of the *cbfunc* callback function serves as notice that the PMIx library no longer requires
14 the caller to maintain the *bo* data object - it does *not* indicate delivery of the payload to the targets.
15 Use of **PMIX_RANK_WILDCARD** to specify all processes in a given namespace is supported but
16 should be used carefully due to bandwidth and memory footprint considerations.

Storage Support Definitions

1 **Provisional** Distributed and parallel computing systems are increasingly embracing storage hierarchies to meet
 2 the diverse data management needs of applications and other systems software in a cost-effective
 3 manner. These hierarchies provide access to a number of distinct storage layers, with each
 4 potentially composed of different storage hardware (e.g., HDD, SSD, tape, PMEM), deployed at
 5 different locations (e.g., on-node, on-switch, on-site, WAN), and designed using different storage
 6 paradigms (e.g., file-based, object-based). Each of these systems offers unique performance and
 7 usage characteristics that storage system users should carefully consider to ensure the most efficient
 8 use of storage resources.

9 PMix enables users to better understand storage hierarchies by defining attributes that formalize
 10 storage system characteristics, state, and other parameters. These attributes can be queried by
 11 applications, I/O libraries and middleware, and workflow systems to discover available storage
 12 resources and to inform on which resources are most suitable for different I/O workload
 13 requirements.

14 19.1 Storage support constants

15 **Provisional** The `pmix_storage_medium_t` is a `uint64_t` type that defines a set of bit-mask flags for
 16 specifying different types of storage mediums. These can be bitwise OR'd together to
 17 accommodate storage systems that mix storage medium types.

18 Provisional	PMIX_STORAGE_MEDIUM_UNKNOWN	0x0000000000000001	The storage medium
19	type is unknown.		
20 Provisional	PMIX_STORAGE_MEDIUM_TAPE	0x0000000000000002	The storage system uses
21	tape media.		
22 Provisional	PMIX_STORAGE_MEDIUM_HDD	0x0000000000000004	The storage system uses
23	HDDs with traditional SAS, SATA interfaces.		
24 Provisional	PMIX_STORAGE_MEDIUM_SSD	0x0000000000000008	The storage system uses
25	SSDs with traditional SAS, SATA interfaces.		
26 Provisional	PMIX_STORAGE_MEDIUM_NVME	0x0000000000000010	The storage system uses
27	SSDs with NVMe interface.		
28 Provisional	PMIX_STORAGE_MEDIUM_PMEM	0x0000000000000020	The storage system uses
29	persistent memory.		
30 Provisional	PMIX_STORAGE_MEDIUM_RAM	0x0000000000000040	The storage system is
31	volatile (e.g., tmpfs).		

Advice to PMIx library implementers

PMIx implementations should maintain the same ordering for bit-mask values for `pmix_storage_medium_t` struct as provided in this standard, since these constants are ordered to provide semantic information that may be of use to PMIx users. Namely, `pmix_storage_medium_t` constants are ordered in terms of increasing medium bandwidth.

It is further recommended that implementations should try to allocate empty bits in the mask so that they can be extended to account for new constant definitions corresponding to new storage mediums.

Provisional

The `pmix_storage_accessibility_t` is a `uint64_t` type that defines a set of bit-mask flags for specifying different levels of storage accessibility (i.e., from where a storage system may be accessed). These can be bitwise OR'd together to accommodate storage systems that are accessible in multiple ways.

Provisional

PMIX_STORAGE_ACCESSIBILITY_NODE 0x0000000000000001 The storage system resources are accessible within the same node.

Provisional

PMIX_STORAGE_ACCESSIBILITY_SESSION 0x0000000000000002 The storage system resources are accessible within the same session.

Provisional

PMIX_STORAGE_ACCESSIBILITY_JOB 0x0000000000000004 The storage system resources are accessible within the same job.

Provisional

PMIX_STORAGE_ACCESSIBILITY_RACK 0x0000000000000008 The storage system resources are accessible within the same rack.

Provisional

PMIX_STORAGE_ACCESSIBILITY_CLUSTER 0x0000000000000010 The storage system resources are accessible within the same cluster.

Provisional

PMIX_STORAGE_ACCESSIBILITY_REMOTE 0x0000000000000020 The storage system resources are remote.

Provisional

The `pmix_storage_persistence_t` is a `uint64_t` type that defines a set of bit-mask flags for specifying different levels of persistence for a particular storage system.

Provisional

PMIX_STORAGE_PERSISTENCE_TEMPORARY 0x0000000000000001 Data on the storage system is persisted only temporarily (i.e., it does not survive across sessions or node reboots).

Provisional

PMIX_STORAGE_PERSISTENCE_NODE 0x0000000000000002 Data on the storage system is persisted on the node.

Provisional

PMIX_STORAGE_PERSISTENCE_SESSION 0x0000000000000004 Data on the storage system is persisted for the duration of the session.

Provisional

PMIX_STORAGE_PERSISTENCE_JOB 0x0000000000000008 Data on the storage system is persisted for the duration of the job.

Provisional

PMIX_STORAGE_PERSISTENCE_SCRATCH 0x0000000000000010 Data on the storage system is persisted according to scratch storage policies (short-term storage, typically persisted for days to weeks).

1	PMIX_STORAGE_PERSISTENCE_PROJECT	0x0000000000000020	Data on the
2			storage system is persisted according to project storage policies (long-term storage, typically
3			persisted for the duration of a project).
4	<i>Provisional</i> PMIX_STORAGE_PERSISTENCE_ARCHIVE	0x0000000000000040	Data on the
5			storage system is persisted according to archive storage policies (long-term storage, typically
6			persisted indefinitely).
7	<i>Provisional</i>		The pmix_storage_access_type_t is a uint16_t type that defines a set of bit-mask
8			flags for specifying different storage system access types.
9	<i>Provisional</i> PMIX_STORAGE_ACCESS_RD	0x0001	Provide information on storage system read
10			operations.
11	<i>Provisional</i> PMIX_STORAGE_ACCESS_WR	0x0002	Provide information on storage system write
12			operations.
13	<i>Provisional</i> PMIX_STORAGE_ACCESS_RDWR	0x0003	Provide information on storage system read
14			and write operations.

15 19.2 Storage support attributes

16 The following attributes may be returned in response to queries (e.g., **PMIx_Get** or
17 **PMIx_Query_info**) made by processes or tools.

18	<i>Provisional</i> PMIX_STORAGE_ID	"pmix.strg.id" (char*)	
19			An identifier for the storage system (e.g., lustre-fs1, daos-oss1, home-fs)
20	<i>Provisional</i> PMIX_STORAGE_PATH	"pmix.strg.path" (char*)	
21			Mount point path for the storage system (valid only for file-based storage systems)
22	<i>Provisional</i> PMIX_STORAGE_TYPE	"pmix.strg.type" (char*)	
23			Type of storage system (i.e., "lustre", "gpfs", "daos", "ext4")
24	<i>Provisional</i> PMIX_STORAGE_VERSION	"pmix.strg.ver" (char*)	
25			Version string for the storage system
26	<i>Provisional</i> PMIX_STORAGE_MEDIUM	"pmix.strg.medium" (pmix_storage_medium_t)	
27			Types of storage mediums utilized by the storage system (e.g., SSDs, HDDs, tape)
28		PMIX_STORAGE_ACCESSIBILITY	
29	<i>Provisional</i>	"pmix.strg.access" (pmix_storage_accessibility_t)	
30			Accessibility level of the storage system (e.g., within same node, within same session)
31		PMIX_STORAGE_PERSISTENCE	
32	<i>Provisional</i>	"pmix.strg.persist" (pmix_storage_persistence_t)	
33			Persistence level of the storage system (e.g., scratch storage or archive storage)
34	<i>Provisional</i> PMIX_QUERY_STORAGE_LIST	"pmix.strg.list" (char*)	
35			Comma-delimited list of storage identifiers (i.e., PMIX_STORAGE_ID types) for available
36			storage systems
37	<i>Provisional</i> PMIX_STORAGE_CAPACITY_LIMIT	"pmix.strg.caplim" (double)	
38			Overall limit on capacity (in bytes) for the storage system
39	<i>Provisional</i> PMIX_STORAGE_CAPACITY_USED	"pmix.strg.capuse" (double)	

1 Overall used capacity (in bytes) for the storage system
2 **PMIX_STORAGE_OBJECT_LIMIT** "pmix.strg.objlim" (uint64_t)
3 Overall limit on number of objects (e.g., inodes) for the storage system
4 *Provisional* **PMIX_STORAGE_OBJECTS_USED** "pmix.strg.objuse" (uint64_t)
5 Overall used number of objects (e.g., inodes) for the storage system
6 *Provisional* **PMIX_STORAGE_MINIMAL_XFER_SIZE** "pmix.strg.minxfer" (double)
7 Minimal transfer size (in bytes) for the storage system - this is the storage system's atomic
8 unit of transfer (e.g., block size)
9 *Provisional* **PMIX_STORAGE_SUGGESTED_XFER_SIZE** "pmix.strg.sxfer" (double)
10 Suggested transfer size (in bytes) for the storage system
11 *Provisional* **PMIX_STORAGE_BW_MAX** "pmix.strg.bwmax" (double)
12 Maximum bandwidth (in bytes/sec) for storage system - provided as the theoretical
13 maximum or the maximum observed bandwidth value
14 *Provisional* **PMIX_STORAGE_BW_CUR** "pmix.strg.bwcur" (double)
15 Observed bandwidth (in bytes/sec) for storage system - provided as a recently observed
16 bandwidth value, with the exact measurement interval depending on the storage system
17 and/or PMIx library implementation
18 *Provisional* **PMIX_STORAGE_IOPS_MAX** "pmix.strg.iopsmax" (double)
19 Maximum IOPS (in I/O operations per second) for storage system - provided as the
20 theoretical maximum or the maximum observed IOPS value
21 *Provisional* **PMIX_STORAGE_IOPS_CUR** "pmix.strg.iopscur" (double)
22 Observed IOPS (in I/O operations per second) for storage system - provided as a recently
23 observed IOPS value, with the exact measurement interval depending on the storage system
24 and/or PMIx library implementation
25 **PMIX_STORAGE_ACCESS_TYPE**
26 *Provisional* "pmix.strg.atype" (pmix_storage_access_type_t)
27 Qualifier describing the type of storage access to return information for (e.g., for qualifying
28 **PMIX_STORAGE_BW_CUR**, **PMIX_STORAGE_IOPS_CUR**, or
29 **PMIX_STORAGE_SUGGESTED_XFER_SIZE** attributes)

APPENDIX A

Python Bindings

1 While the PMIx Standard is defined in terms of C-based APIs, there is no intent to limit the use of
2 PMIx to that specific language. Support for other languages is captured in the Standard by
3 describing their equivalent syntax for the PMIx APIs and native forms for the PMIx datatypes. This
4 Appendix specifically deals with Python interfaces, beginning with a review of the PMIx datatypes.
5 Support is restricted to Python 3 and above - i.e., the Python bindings do not support Python 2.

6 Note: the PMIx APIs have been loosely collected into three Python classes based on their PMIx
7 “class” (i.e., client, server, and tool). All processes have access to a basic set of the APIs, and
8 therefore those have been included in the “client” class. Servers can utilize any of those functions
9 plus a set focused on operations not commonly executed by an application process. Finally, tools
10 can also act as servers but have their own initialization function.

11 A.1 Design Considerations

12 Several issues arose during design of the Python bindings:

13 A.1.1 Error Codes vs Python Exceptions

14 The C programming language reports errors through the return of the corresponding integer status
15 codes. PMIx has defined a range of negative values for this purpose. However, Python has the
16 option of raising *exceptions* that effectively operate as interrupts that can be trapped if the program
17 appropriately tests for them. The PMIx Python bindings opted to follow the C-based standard and
18 return PMIx status codes in lieu of raising exceptions as this method was considered more
19 consistent for those working in both domains.

20 A.1.2 Representation of Structured Data

21 PMIx utilizes a number of C-language structures to efficiently bundle related information. For
22 example, the PMIx process identifier is represented as a struct containing a character array for the
23 namespace and a 32-bit unsigned integer for the process rank. There are several options for
24 translating such objects to Python – e.g., the PMIx process identifier could be represented as a
25 two-element tuple (nspace, rank) or as a dictionary ‘nspace’: name, ‘rank’: 0. Exploration found no
26 discernible benefit to either representation, nor was any clearly identifiable rationale developed that
27 would lead a user to expect one versus the other for a given PMIx data type. Consistency in the
28 translation (i.e., exclusively using tuple or dictionary) appeared to be the most important criterion.
29 Hence, the decision was made to express all complex datatypes as Python dictionaries.

1 A.2 Datatype Definitions

2 PMIx defines a number of datatypes comprised of fixed-size character arrays, restricted range
3 integers (e.g., `uint32_t`), and structures. Each datatype is represented by a named unsigned 16-bit
4 integer (`uint16_t`) constant. Users are advised to use the named PMIx constants for indicating
5 datatypes instead of integer values to ensure compatibility with future PMIx versions.

6 With only a few exceptions, the C-based PMIx datatypes defined in Chapter 3 on page 13 directly
7 translate to Python. However, Python lacks the size-specific value definitions of C (e.g., `uint8_t`)
8 and thus some care must be taken to protect against overflow/underflow situations when moving
9 between the languages. Python bindings that accept values including PMIx datatypes shall
10 therefore have the datatype and associated value checked for compatibility with their PMIx-defined
11 equivalents, returning an error if:

- 12 • datatypes not defined by PMIx are encountered
- 13 • provided values fall outside the range of the C-equivalent definition - e.g., if a value identified as
14 `PMIX_UINT8` lies outside the `uint8_t` range

15 Note that explicit labeling of PMIx data type, even when Python itself doesn't care, is often
16 required for the Python bindings to know how to properly interpret and label the provided value
17 when passing it to the PMIx library.

18 Table A.1 lists the correspondence between data types in the two languages.

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>bool</code>	<code>PMIX_BOOL</code>	boolean	
<code>byte</code>	<code>PMIX_BYTE</code>	A single element byte array (i.e., a byte array of length one)	
<code>char*</code>	<code>PMIX_STRING</code>	string	
<code>size_t</code>	<code>PMIX_SIZE</code>	integer	
<code>pid_t</code>	<code>PMIX_PID</code>	integer	value shall be limited to the <code>uint32_t</code> range
<code>int, int8_t, int16_t, int32_t, int64_t</code>	<code>PMIX_INT, PMIX_INT8, PMIX_INT16, PMIX_INT32, PMIX_INT64</code>	integer	value shall be limited to its corresponding range
<code>uint, uint8_t, uint16_t, uint32_t, uint64_t</code>	<code>PMIX_UINT, PMIX_UINT8, PMIX_UINT16, PMIX_UINT32, PMIX_UINT64</code>	integer	value shall be limited to its corresponding range
<code>float, double</code>	<code>PMIX_FLOAT, PMIX_DOUBLE</code>	float	value shall be limited to its corresponding range
<code>struct timeval</code>	<code>PMIX_TIMEVAL</code>	{ 'sec': sec, 'usec': microsec }	each field is an integer value
<code>time_t</code>	<code>PMIX_TIME</code>	integer	limited to positive values
<code>pmix_data_type_t</code>	<code>PMIX_DATA_TYPE</code>	integer	value shall be limited to the <code>uint16_t</code> range
<code>pmix_status_t</code>	<code>PMIX_STATUS</code>	integer	
<code>pmix_key_t</code>	N/A	string	The string's length shall be limited to one less than the size of the <code>pmix_key_t</code> array (to reserve space for the terminating NULL)
<code>pmix_nspace_t</code>	N/A	string	The string's length shall be limited to one less than the size of the <code>pmix_nspace_t</code> array (to reserve space for the terminating NULL)

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>pmix_rank_t</code>	<code>PMIX_PROC_RANK</code>	integer	value shall be limited to the <code>uint32_t</code> range excepting the reserved values near <code>UINT32_MAX</code>
<code>pmix_proc_t</code>	<code>PMIX_PROC</code>	{'nspace': nspace, 'rank': rank}	<i>nspace</i> is a Python string and <i>rank</i> is an integer value. The <i>nspace</i> string's length shall be limited to one less than the size of the <code>pmix_nspace_t</code> array (to reserve space for the terminating <code>NULL</code>), and the <i>rank</i> value shall conform to the constraints associated with <code>pmix_rank_t</code>
<code>pmix_byte_object_t</code>	<code>PMIX_BYTE_OBJECT</code>	{'bytes': bytes, 'size': size}	<i>bytes</i> is a Python byte array and <i>size</i> is the integer number of bytes in that array.
<code>pmix_persistence_t</code>	<code>PMIX_PERSISTENCE</code>	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_scope_t</code>	<code>PMIX_SCOPE</code>	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_data_range_t</code>	<code>PMIX_RANGE</code>	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_proc_state_t</code>	<code>PMIX_PROC_STATE</code>	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_proc_info_t</code>	<code>PMIX_PROC_INFO</code>	{'proc': {'nspace': nspace, 'rank': rank}, 'hostname': hostname, 'executable': executable, 'pid': pid, 'exitcode': exitcode, 'state': state}	<i>proc</i> is a Python <code>proc</code> dictionary; <i>hostname</i> and <i>executable</i> are Python strings; and <i>pid</i> , <i>exitcode</i> , and <i>state</i> are Python integers

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>pmix_data_array_t</code>	<code>PMIX_DATA_ARRAY</code>	{'type': type, 'array': array}	<i>type</i> is the PMIx type of object in the array and <i>array</i> is a Python <i>list</i> containing the individual array elements. Note that <i>array</i> can consist of <i>any</i> PMIx types, including (for example) a Python <code>info</code> object that itself contains an <code>array</code> value
<code>pmix_info_directives_t</code>	<code>PMIX_INFO_DIRECTIVES</code>	list	list of integer values (defined in Section 3.2.10)
<code>pmix_alloc_directive_t</code>	<code>PMIX_ALLOC_DIRECTIVE</code>	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_iof_channel_t</code>	<code>PMIX_IOF_CHANNEL</code>	list	list of integer values (defined in Section 18.3.3)
<code>pmix_envvar_t</code>	<code>PMIX_ENVAR</code>	{'envar': envar, 'value': value, 'separator': separator}	<i>envar</i> and <i>value</i> are Python strings, and <i>separator</i> a single-character Python string
<code>pmix_value_t</code>	<code>PMIX_VALUE</code>	{'value': value, 'val_type': type}	<i>type</i> is the PMIx datatype of <i>value</i> , and <i>value</i> is the associated value expressed in the appropriate Python form for the specified datatype
<code>pmix_info_t</code>	<code>PMIX_INFO</code>	{'key': key, 'flags': flags, 'value': value, 'val_type': type}	<i>key</i> is a Python string <code>key</code> , <i>flags</i> is an <code>info directives</code> value, <i>type</i> is the PMIx datatype of <i>value</i> , and <i>value</i> is the associated value expressed in the appropriate Python form for the specified datatype
<code>pmix_pdata_t</code>	<code>PMIX_PDATA</code>	{'proc': {'nspace': nspace, 'rank': rank}, 'key': key, 'value': value, 'val_type': type}	<i>proc</i> is a Python <code>proc</code> dictionary; <i>key</i> is a Python string <code>key</code> ; <i>type</i> is the PMIx datatype of <i>value</i> ; and <i>value</i> is the associated value expressed in the appropriate Python form for the specified datatype

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>pmix_app_t</code>	<code>PMIX_APP</code>	{'cmd': cmd, 'argv': [argv], 'env': [env], 'maxprocs': maxprocs, 'info': [info]}	<i>cmd</i> is a Python string; <i>argv</i> and <i>env</i> are Python <i>lists</i> containing Python strings; <i>maxprocs</i> is an integer; and <i>info</i> is a Python <i>list</i> of info values
<code>pmix_query_t</code>	<code>PMIX_QUERY</code>	{'keys': [keys], 'qualifiers': [info]}	<i>keys</i> is a Python <i>list</i> of Python strings, and <i>qualifiers</i> is a Python <i>list</i> of info values
<code>pmix_regattr_t</code>	<code>PMIX_REGATTR</code>	{'name': name, 'key': key, 'type': type, 'info': [info], 'description': [desc]}	<i>name</i> and <i>string</i> are Python strings; <i>type</i> is the PMIx datatype for the attribute's value; <i>info</i> is a Python <i>list</i> of info values; and <i>description</i> is a list of Python strings describing the attribute
<code>pmix_job_state_t</code>	<code>PMIX_JOB_STATE</code>	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_link_state_t</code>	<code>PMIX_LINK_STATE</code>	integer	value shall be limited to the <code>uint8_t</code> range
<code>pmix_cpuset_t</code>	<code>PMIX_PROC_CPuset</code>	{'source': source, 'cpus': bitmap}	<i>source</i> is a string name of the library that created the cpuset; and <i>cpus</i> is a list of string ranges identifying the PUs to which the process is bound (e.g., [1, 3-5, 7])
<code>pmix_locality_t</code>	<code>PMIX_LOCTYPE</code>	list	list of integer values (defined in Section 12.4.2.3) describing the relative locality of the specified local process
<code>pmix_fabric_t</code>	N/A	{'name': name, 'index': idx, 'info': [info]}	<i>name</i> is the string name assigned to the fabric; <i>index</i> is the integer ID assigned to the fabric; <i>info</i> is a list of info describing the fabric
<code>pmix_endpoint_t</code>	<code>PMIX_ENDPOINT</code>	{'uuid': uuid, 'osname': osname, 'endpt': endpt}	<i>uuid</i> is the string system-unique identifier assigned to the device; <i>osname</i> is the operating system name assigned to the device; <i>endpt</i> is a byteobject containing the endpoint information

Table A.1.: C-to-Python Datatype Correspondence

C-Definition	PMIx Name	Python Definition	Notes
<code>pmix_device_distance_t</code>	<code>PMIX_DEVICE_DIST</code>	{'uuid': uuid, 'osname': osname, 'mindist': mindist, 'maxdist': maxdist}	<i>uuid</i> is the string system-unique identifier assigned to the device; <i>osname</i> is the operating system name assigned to the device; and <i>mindist</i> and <i>maxdist</i> are Python integers
<code>pmix_coord_t</code>	<code>PMIX_COORD</code>	{'view': view, 'coord': [coords]}	<i>view</i> is the <code>pmix_coord_view_t</code> of the coordinate; and <i>coord</i> is a list of integer coordinates, one for each dimension of the fabric
<code>pmix_geometry_t</code>	<code>PMIX_GEOMETRY</code>	{'fabric': idx, 'uuid': uuid, 'osname': osname, 'coordinates': [coords]}	<i>fabric</i> is the Python integer index of the fabric; <i>uuid</i> is the string system-unique identifier assigned to the device; <i>osname</i> is the operating system name assigned to the device; and <i>coordinates</i> is a list of <code>coord</code> containing the coordinates for the device across all views
<code>pmix_device_type_t</code>	<code>PMIX_DEVTYPE</code>	list	list of integer values (defined in Section 12.4.8)
<code>pmix_bind_envelope_t</code>	N/A	integer	one of the values defined in Section 12.4.4.1

1 A.3 Callback Function Definitions

2 A.3.1 IOF Delivery Function

3 Summary

4 Callback function for delivering forwarded IO to a process

5 *PMIx v4.0*

Format

Python

```
6 def iofcbfunc(iofhdlr:integer, channel:bitarray,  
7             source:dict, payload:dict, info:list)
```

Python

8 **IN** `iofhdlr`

Registration number of the handler being invoked (integer)

9 **IN** `channel`

Python `channel` 16-bit bitarray identifying the channel the data arrived on (bitarray)

10 **IN** `source`

Python `proc` identifying the namespace/rank of the process that generated the data (dict)

11 **IN** `payload`

Python `byteobject` containing the data (dict)

12 **IN** `info`

List of Python `info` provided by the source containing metadata about the payload. This could include `PMIX_IOF_COMPLETE` (list)

13 Returns: nothing

14 See `pmix_iof_cbfunc_t` for details

21 A.3.2 Event Handler

22 Summary

23 Callback function for event handlers

24 *PMIx v4.0* Format

Python

```
1 def evhandler(evhdlr:integer, status:integer,  
2             source:dict, info:list, results:list)
```

Python

3 **IN** `iofhdlr`

4 Registration number of the handler being invoked (integer)

5 **IN** `status`

6 Status associated with the operation (integer)

7 **IN** `source`

8 Python `proc` identifying the namespace/rank of the process that generated the event (dict)

9 **IN** `info`

10 List of Python `info` provided by the source containing metadata about the event (list)

11 **IN** `results`

12 List of Python `info` containing the aggregated results of all prior evhandlers (list)

13 Returns:

- 14 • `rc` - Status returned by the event handler's operation (integer)
- 15 • `results` - List of Python `info` containing results from this event handler's operation on the event
- 16 (list)

17 See `pmix_notification_fn_t` for details

18 A.3.3 Server Module Functions

19 The following definitions represent functions that may be provided to the PMIx server library at
20 time of initialization for servicing of client requests. Module functions that are not provided default
21 to returning "not supported" to the caller.

22 A.3.3.1 Client Connected

23 Summary

24 Notify the host server that a client connected to this server.

25 Format

PMIx v4.0

Python

1 `def clientconnected2(proc:dict is not None, info:list)`

Python

2 **IN** `proc`

3 Python `proc` identifying the namespace/rank of the process that connected (dict)

4 **IN** `info`

5 list of Python `info` containing information about the process (list)

6 Returns:

7 • `rc` - `PMIX_SUCCESS` or a PMIx error code indicating the connection should be rejected (integer)

8 See `pmix_server_client_connected2_fn_t` for details

9 **A.3.3.2 Client Finalized**

10 **Summary**

11 Notify the host environment that a client called `PMIx_Finalize`.

12 **Format**

PMIx v4.0

Python

13 `def clientfinalized(proc:dict is not None):`

Python

14 **IN** `proc`

15 Python `proc` identifying the namespace/rank of the process that finalized (dict)

16 Returns: nothing

17 See `pmix_server_client_finalized_fn_t` for details

18 **A.3.3.3 Client Aborted**

19 **Summary**

20 Notify the host environment that a local client called `PMIx_Abort`.

Format

Python

```
def clientaborted(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'caller': Python **proc** identifying the namespace/rank of the process calling abort (dict)
- 'status': PMIx status to be returned on exit (integer)
- 'msg': Optional string message to be printed (string)
- 'targets': Optional list of Python **proc** identifying the namespace/rank of the processes to be aborted (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

See [pmix_server_abort_fn_t](#) for details

A.3.3.4 Fence

Summary

At least one client called either [PMIx_Fence](#) or [PMIx_Fence_nb](#)

Format

Python

PMIx v4.0

```
def fence(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'procs': List of Python **proc** identifying the namespace/rank of the participating processes (list)
- 'directives': Optional list of Python **info** containing directives controlling the operation (list)
- 'data': Optional Python bytearray of data to be circulated during fence operation (bytearray)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- *data* - Python bytearray containing the aggregated data from all participants (bytearray)

See [pmix_server_fence_nb_fn_t](#) for details

1 **A.3.3.5 Direct Modex**

2 **Summary**

3 Used by the PMIx server to request its local host contact the PMIx server on the remote node that
4 hosts the specified proc to obtain and return a direct modex blob for that proc.

5 **Format**

PMIx v4.0

Python

```
6 def dmodex(args:dict is not None)
```

Python

7 **IN args**

8 Python dictionary containing:

- 9 • 'proc': Python **proc** of process whose data is being requested (dict)
- 10 • 'directives': Optional list of Python **info** containing directives controlling the operation
11 (list)

12 Returns:

- 13 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- 14 • *data* - Python bytearray containing the data for the specified process (bytearray)

15 See [pmix_server_dmodex_req_fn_t](#) for details

16 **A.3.3.6 Publish**

17 **Summary**

18 Publish data per the PMIx API specification.

19 **Format**

PMIx v4.0

Python

```
20 def publish(args:dict is not None)
```

Python

21 **IN args**

22 Python dictionary containing:

- 23 • 'proc': Python **proc** dictionary of process publishing the data (dict)
- 24 • 'directives': List of Python **info** containing data and directives (list)

25 Returns:

- 26 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

27 See [pmix_server_publish_fn_t](#) for details

1 **A.3.3.7 Lookup**

2 **Summary**

3 Lookup published data.

4 **Format**

PMIx v4.0

Python

```
5 def lookup(args:dict is not None)
```

Python

6 **IN args**

7 Python dictionary containing:

- 8 • 'proc': Python **proc** of process seeking the data (dict)
- 9 • 'keys': List of Python strings (list)
- 10 • 'directives': Optional list of Python **info** containing directives (list)

11 Returns:

- 12 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- 13 • *pdata* - List of **pdata** containing the returned results (list)

14 See [pmix_server_lookup_fn_t](#) for details

15 **A.3.3.8 Unpublish**

16 **Summary**

17 Delete data from the data store.

18 **Format**

PMIx v4.0

Python

```
19 def unpublish(args:dict is not None)
```

Python

20 **IN args**

21 Python dictionary containing:

- 22 • 'proc': Python **proc** of process unpublishing data (dict)
- 23 • 'keys': List of Python strings (list)
- 24 • 'directives': Optional list of Python **info** containing directives (list)

25 Returns:

- 26 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

27 See [pmix_server_unpublish_fn_t](#) for details

1 **A.3.3.9 Spawn**

2 **Summary**

3 Spawn a set of applications/processes as per the [PMIx_Spawn](#) API.

4 *PMIx v4.0* **Format**

Python

```
5 def spawn(args:dict is not None)
```

Python

6 **IN args**

7 Python dictionary containing:

- 8 • 'proc': Python [proc](#) of process making the request (dict)
- 9 • 'jobinfo': Optional list of Python [info](#) job-level directives and information (list)
- 10 • 'apps': List of Python [app](#) describing applications to be spawned (list)

11 Returns:

- 12 • *rc* - [PMIX_SUCCESS](#) or a PMIx error code indicating the operation failed (integer)
- 13 • *nspace* - Python string containing namespace of the spawned job (str)

14 See [pmix_server_spawn_fn_t](#) for details

15 **A.3.3.10 Connect**

16 **Summary**

17 Record the specified processes as *connected*.

18 *PMIx v4.0* **Format**

Python

```
19 def connect(args:dict is not None)
```

Python

20 **IN args**

21 Python dictionary containing:

- 22 • 'procs': List of Python [proc](#) identifying the namespace/rank of the participating processes (list)
- 23 • 'directives': Optional list of Python [info](#) containing directives controlling the operation (list)

26 Returns:

- 27 • *rc* - [PMIX_SUCCESS](#) or a PMIx error code indicating the operation failed (integer)

28 See [pmix_server_connect_fn_t](#) for details

1 **A.3.3.11 Disconnect**

2 **Summary**

3 Disconnect a previously connected set of processes.

4 **Format**

PMIx v4.0

Python

```
5 def disconnect(args:dict is not None)
```

Python

6 **IN args**

7 Python dictionary containing:

- 8 • 'procs': List of Python **proc** identifying the namespace/rank of the participating processes (list)
- 9 • 'directives': Optional list of Python **info** containing directives controlling the operation (list)

12 Returns:

- 13 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

14 See [pmix_server_disconnect_fn_t](#) for details

15 **A.3.3.12 Register Events**

16 **Summary**

17 Register to receive notifications for the specified events.

18 **Format**

PMIx v4.0

Python

```
19 def register_events(args:dict is not None)
```

Python

20 **IN args**

21 Python dictionary containing:

- 22 • 'codes': List of Python integers (list)
- 23 • 'directives': Optional list of Python **info** containing directives controlling the operation (list)

25 Returns:

- 26 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

27 See [pmix_server_register_events_fn_t](#) for details

1 A.3.3.13 Deregister Events

2 Summary

3 Deregister to receive notifications for the specified events.

4 *PMIx v4.0* Format

Python

```
5 def deregister_events(args:dict is not None)
```

Python

6 IN args

7 Python dictionary containing:

- 8 • 'codes': List of Python integers (list)

9 Returns:

- 10 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

11 See [pmix_server_deregister_events_fn_t](#) for details

12 A.3.3.14 Notify Event

13 Summary

14 Notify the specified range of processes of an event.

15 *PMIx v4.0* Format

Python

```
16 def notify_event(args:dict is not None)
```

Python

17 IN args

18 Python dictionary containing:

- 19 • 'code': Python integer [pmix_status_t](#) (integer)
- 20 • 'source': Python [proc](#) of process that generated the event (dict)
- 21 • 'range': Python [range](#) in which the event is to be reported (integer)
- 22 • 'directives': Optional list of Python [info](#) directives (list)

23 Returns:

- 24 • *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

25 See [pmix_server_notify_event_fn_t](#) for details

26 A.3.3.15 Query

27 Summary

28 Query information from the resource manager.

1 **Format** Python

2 `def query(args:dict is not None)`
3 Python

4 **IN** `args`
5 Python dictionary containing:
6 • 'source': Python **proc** of requesting process (dict)
7 • 'queries': List of Python **query** directives (list)

8 Returns:
9 • `rc` - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
10 • `info` - List of Python **info** containing the returned results (list)

11 See [pmix_server_query_fn_t](#) for details

11 A.3.3.16 Tool Connected

12 **Summary**
13 Register that a tool has connected to the server.

14 *PMIx v4.0* **Format** Python

15 `def tool_connected(args:dict is not None)`
16 Python

17 **IN** `args`
18 Python dictionary containing:
19 • 'directives': Optional list of Python **info** info on the connecting tool (list)

20 Returns:
21 • `rc` - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
22 • `proc` - Python **proc** containing the assigned namespace:rank for the tool (dict)

23 See [pmix_server_tool_connection_fn_t](#) for details

23 A.3.3.17 Log

24 **Summary**
25 Log data on behalf of a client.

Format

Python

```
def log(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'source': Python **proc** of requesting process (dict)
- 'data': Optional list of Python **info** containing data to be logged (list)
- 'directives': Optional list of Python **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

See [pmix_server_log_fn_t](#) for details.

A.3.3.18 Allocate Resources

Summary

Request allocation operations on behalf of a client.

Format

Python

```
def allocate(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'source': Python **proc** of requesting process (dict)
- 'action': Python **allocdir** specifying requested action (integer)
- 'directives': Optional list of Python **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- *refarginfo* - List of Python **info** containing results of requested operation (list)

See [pmix_server_alloc_fn_t](#) for details.

A.3.3.19 Job Control

Summary

Execute a job control action on behalf of a client.

Format

Python

```
def job_control(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'source': Python **proc** of requesting process (dict)
- 'targets': List of Python **proc** specifying target processes (list)
- 'directives': Optional list of Python **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

See [pmix_server_job_control_fn_t](#) for details.

A.3.3.20 Monitor

Summary

Request that a client be monitored for activity.

Format

Python

```
def monitor(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'source': Python **proc** of requesting process (dict)
- 'monitor': Python **info** attribute indicating the type of monitor being requested (dict)
- 'error': Status code to be used when generating an event notification (integer) alerting that the monitor has been triggered.
- 'directives': Optional list of Python **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)

See [pmix_server_monitor_fn_t](#) for details.

A.3.3.21 Get Credential

Summary

Request a credential from the host environment.

Format

Python

```
def get_credential(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'source': Python **proc** of requesting process (dict)
- 'directives': Optional list of Python **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- *cred* - Python **byteobject** containing returned credential (dict)
- *info* - List of Python **info** containing any additional info about the credential (list)

See [pmix_server_get_cred_fn_t](#) for details.

A.3.3.22 Validate Credential

Summary

Request validation of a credential

Format

Python

```
def validate_credential(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'source': Python **proc** of requesting process (dict)
- 'credential': Python **byteobject** containing credential (dict)
- 'directives': Optional list of Python **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
- *info* - List of Python **info** containing any additional info from the credential (list)

See [pmix_server_validate_cred_fn_t](#) for details.

A.3.3.23 IO Forward

Summary

Request the specified IO channels be forwarded from the given array of processes.

Format

Python

```
def iof_pull(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'sources': List of Python [proc](#) of processes whose IO is being requested (list)
- 'channels': Bitmask of Python [channel](#) identifying IO channels to be forwarded (integer)
- 'directives': Optional list of Python [info](#) containing directives (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a PMIx error code indicating the operation failed (integer)

See [pmix_server_iof_fn_t](#) for details.

A.3.3.24 IO Push

Summary

Pass standard input data to the host environment for transmission to specified recipients.

Format

Python

```
def iof_push(args:dict is not None)
```

Python

IN args

Python dictionary containing:

- 'source': Python [proc](#) of process whose input is being forwarded (dict)
- 'payload': Python [byteobject](#) containing input bytes (dict)
- 'targets': List of [proc](#) of processes that are to receive the payload (list)
- 'directives': Optional list of Python [info](#) containing directives (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a PMIx error code indicating the operation failed (integer)

See [pmix_server_stdin_fn_t](#) for details.

A.3.3.25 Group Operations

Summary

Request group operations (construct, destruct, etc.) on behalf of a set of processes.

1 **Format** Python

2 `def group(args:dict is not None)` Python

3 **IN args**
4 Python dictionary containing:
5 • 'op': Operation host is to perform on the specified group (integer)
6 • 'group': String identifier of target group (str)
7 • 'procs': List of Python **proc** of participating processes (dict)
8 • 'directives': Optional list of Python **info** containing directives (list)

9 Returns:
10 • rc - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
11 • refarginfo - List of Python **info** containing results of requested operation (list)
12 See [pmix_server_grp_fn_t](#) for details.

13 **A.3.3.26 Fabric Operations**

14 **Summary**
15 Request fabric-related operations (e.g., information on a fabric) on behalf of a tool or other process.

16 *PMIx v4.0* **Format** Python

17 `def fabric(args:dict is not None)` Python

18 **IN args**
19 Python dictionary containing:
20 • 'source': Python **proc** of requesting process (dict)
21 • 'index': Identifier of the fabric being operated upon (integer)
22 • 'op': Operation host is to perform on the specified fabric (integer)
23 • 'directives': Optional list of Python **info** containing directives (list)

24 Returns:
25 • rc - **PMIX_SUCCESS** or a PMIx error code indicating the operation failed (integer)
26 • refarginfo - List of Python **info** containing results of requested operation (list)
27 See [pmix_server_fabric_fn_t](#) for details.

1 A.4 PMIxClient

2 The client Python class is by far the richest in terms of APIs as it houses all the APIs that an
3 application might utilize. Due to the datatype translation requirements of the C-Python interface,
4 only the blocking form of each API is supported – providing a Python callback function directly to
5 the C interface underlying the bindings was not a supportable option.

6 A.4.1 Client.init

7 Summary

8 Initialize the PMIx client library after obtaining a new PMIxClient object.

9 *PMIx v4.0* Format

Python

```
10 rc, proc = myclient.init(info:list)
```

Python

11 IN info

12 List of Python **info** dictionaries (list)

13 Returns:

- 14 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 15 • *proc* - a Python **proc** dictionary (dict)

16 See **PMIx_Init** for description of all relevant attributes and behaviors.

17 A.4.2 Client.initialized

18 *PMIx v4.0* Format

Python

```
19 rc = myclient.initialized()
```

Python

20 Returns:

- 21 • *rc* - a value of **1** (true) will be returned if the PMIx library has been initialized, and **0** (false)
22 otherwise (integer)

23 See **PMIx_Initialized** for description of all relevant attributes and behaviors.

1 A.4.3 Client.get_version

2 **Format** Python

3 `vers = myclient.get_version()`
4 Python

4 Returns:

- 5 • *vers* - Python string containing the version of the PMIx library (e.g., "3.1.4") (integer)

6 See [PMIx_Get_version](#) for description of all relevant attributes and behaviors.

7 A.4.4 Client.finalize

8 **Summary**

9 Finalize the PMIx client library.

10 *PMIx v4.0* **Format** Python

11 `rc = myclient.finalize(info:list)`
12 Python

12 **IN** *info*

13 List of Python *info* dictionaries (list)

14 Returns:

- 15 • *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

16 See [PMIx_Finalize](#) for description of all relevant attributes and behaviors.

17 A.4.5 Client.abort

18 **Summary**

19 Request that the provided list of processes be aborted.

Format

Python

```
rc = myclient.abort(status:integer, msg:str, targets:list)
```

Python

IN status

PMIx status to be returned on exit (integer)

IN msg

String message to be printed (string)

IN targets

List of Python `proc` dictionaries (list)

Returns:

- `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

See `PMIx_Abort` for description of all relevant attributes and behaviors.

A.4.6 Client.store_internal

Summary

Store some data locally for retrieval by other areas of the process

Format

Python

PMIx v4.0

```
rc = myclient.store_internal(proc:dict, key:str, value:dict)
```

Python

IN proc

Python `proc` dictionary of the process being referenced (dict)

IN key

String key of the data (string)

IN value

Python `value` dictionary (dict)

Returns:

- `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

See `PMIx_Store_internal` for details.

A.4.7 Client.put

Summary

Push a key/value pair into the client's namespace.

1

Format

Python

2

```
rc = myclient.put(scope:integer, key:str, value:dict)
```

Python

3

IN scope

Scope of the data being posted (integer)

4

IN key

String key of the data (string)

5

6

IN value

Python **value** dictionary (dict)

7

8

9

Returns:

10

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

11

See **PMIx_Put** for description of all relevant attributes and behaviors.

12

A.4.8 Client.commit

13

Summary

14

Push all previously **PMIxClient.put** values to the local PMIx server.

15

PMIx v4.0

Format

Python

16

```
rc = myclient.commit()
```

Python

17

Returns:

18

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

19

See **PMIx_Commit** for description of all relevant attributes and behaviors.

20

A.4.9 Client.fence

21

Summary

22

Execute a blocking barrier across the processes identified in the specified list.

Format

Python

```
rc = myclient.fence(peers:list, directives:list)
```

Python

IN peers

List of Python **proc** dictionaries (list)

IN directives

List of Python **info** dictionaries (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See **PMIx_Fence** for description of all relevant attributes and behaviors.

A.4.10 Client.get

Summary

Retrieve a key/value pair.

Format

Python

PMIx v4.0

```
rc, val = myclient.get(proc:dict, key:str, directives:list)
```

Python

IN proc

Python **proc** whose data is being requested (dict)

IN key

Python string key of the data to be returned (str)

IN directives

List of Python **info** dictionaries (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *val* - Python **value** containing the returned data (dict)

See **PMIx_Get** for description of all relevant attributes and behaviors.

A.4.11 Client.publish

Summary

Publish data for later access via **PMIx_Lookup**.

1 **Format** Python

2 `rc = myclient.publish(directives:list)` Python

3 **IN directives**
4 List of Python **info** dictionaries containing data to be published and directives (list)

5 Returns:
6 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
7 See **PMIx_Publish** for description of all relevant attributes and behaviors.

8 **A.4.12 Client.lookup**

9 **Summary**
10 Lookup information published by this or another process with **PMIx_Publish**.

11 **Format** Python

12 `rc, info = myclient.lookup(pdata:list, directives:list)` Python

13 **IN pdata**
14 List of Python **pdata** dictionaries identifying data to be retrieved (list)

15 **IN directives**
16 List of Python **info** dictionaries (list)

17 Returns:
18 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
19 • *info* - Python list of **info** containing the returned data (list)
20 See **PMIx_Lookup** for description of all relevant attributes and behaviors.

21 **A.4.13 Client.unpublish**

22 **Summary**
23 Delete data published by this process with **PMIx_Publish**.

Format

Python

```
rc = myclient.unpublish(keys:list, directives:list)
```

Python

IN keys

List of Python string keys identifying data to be deleted (list)

IN directives

List of Python [info](#) dictionaries (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_Unpublish](#) for description of all relevant attributes and behaviors.

A.4.14 Client.spawn

Summary

Spawn a new job.

Format

Python

PMIx v4.0

```
rc, nspace = myclient.spawn(jobinfo:list, apps:list)
```

Python

IN jobinfo

List of Python [info](#) dictionaries (list)

IN apps

List of Python [app](#) dictionaries (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

- *nspace* - Python [nspace](#) of the new job (dict)

See [PMIx_Spawn](#) for description of all relevant attributes and behaviors.

A.4.15 Client.connect

Summary

Connect namespaces.

1
2
3
4
5
6
7
8
9

Format

Python

```
rc = myclient.connect(peers:list, directives:list)
```

Python

IN peers

List of Python **proc** dictionaries (list)

IN directives

List of Python **info** dictionaries (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- See **PMIx_Connect** for description of all relevant attributes and behaviors.

A.4.16 Client.disconnect

Summary

Disconnect namespaces.

13 *PMIx v4.0*

Format

Python

```
rc = myclient.disconnect(peers:list, directives:list)
```

Python

IN peers

List of Python **proc** dictionaries (list)

IN directives

List of Python **info** dictionaries (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- See **PMIx_Disconnect** for description of all relevant attributes and behaviors.

A.4.17 Client.resolve_peers

Summary

Return list of processes within the specified **nspc** on the given node.

23
24

Format

Python

```
rc,procs = myclient.resolve_peers(node:str, nspace:str)
```

Python

IN node

Name of node whose processes are being requested (str)

IN nspace

Python **nnamespace** whose processes are to be returned (str)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *procs* - List of Python **proc** dictionaries (list)

See **PMIx_Resolve_peers** for description of all relevant attributes and behaviors.

A.4.18 Client.resolve_nodes

Summary

Return list of nodes hosting processes within the specified **nnamespace**.

Format

Python

```
rc,nodes = myclient.resolve_nodes(nspace:str)
```

Python

IN nspace

Python **nnamespace** (str)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *nodes* - List of Python string node names (list)

See **PMIx_Resolve_nodes** for description of all relevant attributes and behaviors.

A.4.19 Client.query

Summary

Query information about the system in general.

1
2
3
4
5
6
7
8

Format

Python

```
rc, info = myclient.query(queries:list)
```

Python

IN queries

List of Python **query** dictionaries (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python **info** containing results of the query (list)

See **PMIx_Query_info** for description of all relevant attributes and behaviors.

9 **A.4.20 Client.log**

Summary

Log data to a central data service/store.

12 *PMIx v4.0*

Format

Python

```
rc = myclient.log(data:list, directives:list)
```

Python

IN data

List of Python **info** (list)

IN directives

Optional list of Python **info** (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See **PMIx_Log** for description of all relevant attributes and behaviors.

21 **A.4.21 Client.allocation_request**

Summary

Request an allocation operation from the host resource manager.

22
23

Format

Python

```
rc, info = myclient.allocation_request(request:integer, directives:list)
```

Python

IN request

Python [allocdir](#) specifying requested operation (integer)

IN directives

List of Python [info](#) describing request (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python [info](#) containing results of the request (list)

See [PMIx_Allocation_request](#) for description of all relevant attributes and behaviors.

A.4.22 Client.job_ctrl

Summary

Request a job control action.

Format

Python

```
rc, info = myclient.job_ctrl(targets:list, directives:list)
```

Python

IN targets

List of Python [proc](#) specifying targets of requested operation (integer)

IN directives

List of Python [info](#) describing operation to be performed (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python [info](#) containing results of the request (list)

See [PMIx_Job_control](#) for description of all relevant attributes and behaviors.

A.4.23 Client.monitor

Summary

Request that something be monitored.

Format

Python

```
rc, info = myclient.monitor(monitor:dict, error_code:integer, directives:list)
```

Python

IN monitor

Python [info](#) specifying specifying the type of monitor being requested (dict)

IN error_code

Status code to be used when generating an event notification alerting that the monitor has been triggered (integer)

IN directives

List of Python [info](#) describing request (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python [info](#) containing results of the request (list)

See [PMIx_Process_monitor](#) for description of all relevant attributes and behaviors.

A.4.24 Client.get_credential

Summary

Request a credential from the PMIx server/SMS.

Format

Python

```
rc, cred = myclient.get_credential(directives:list)
```

Python

IN directives

Optional list of Python [info](#) describing request (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *cred* - Python [byteobject](#) containing returned credential (dict)

See [PMIx_Get_credential](#) for description of all relevant attributes and behaviors.

A.4.25 Client.validate_credential

Summary

Request validation of a credential by the PMIx server/SMS.

Format

Python

```
rc, info = myclient.validate_credential(cred:dict, directives:list)
```

Python

IN cred

Python [byteobject](#) containing credential (dict)

IN directives

Optional list of Python [info](#) describing request (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python [info](#) containing additional results of the request (list)

See [PMIx_Validate_credential](#) for description of all relevant attributes and behaviors.

A.4.26 Client.group_construct

Summary

Construct a new group composed of the specified processes and identified with the provided group identifier.

Format

Python

```
rc, info = myclient.construct_group(grp:string,  
                                   members:list, directives:list)
```

Python

IN grp

Python string identifier for the group (str)

IN members

List of Python [proc](#) dictionaries identifying group members (list)

IN directives

Optional list of Python [info](#) describing request (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python [info](#) containing results of the request (list)

See [PMIx_Group_construct](#) for description of all relevant attributes and behaviors.

A.4.27 Client.group_invite

Summary

Explicitly invite specified processes to join a group.

Format

Python

```
rc, info = myclient.group_invite(grp:string,  
                                members:list, directives:list)
```

Python

IN grp

Python string identifier for the group (str)

IN members

List of Python **proc** dictionaries identifying processes to be invited (list)

IN directives

Optional list of Python **info** describing request (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python **info** containing results of the request (list)

See **PMIx_Group_invite** for description of all relevant attributes and behaviors.

A.4.28 Client.group_join

Summary

Respond to an invitation to join a group that is being asynchronously constructed.

Format

Python

```
rc, info = myclient.group_join(grp:string,  
                               leader:dict, opt:integer,  
                               directives:list)
```

Python

IN grp

Python string identifier for the group (str)

IN leader

Python **proc** dictionary identifying process leading the group (dict)

IN opt

One of the **pmix_group_opt_t** values indicating decline/accept (integer)

IN directives

Optional list of Python **info** describing request (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *info* - List of Python **info** containing results of the request (list)

See **PMIx_Group_join** for description of all relevant attributes and behaviors.

PMIx v4.0

1 A.4.29 Client.group_leave

2 Summary

3 Leave a PMIx Group.

4 *PMIx v4.0* Format

Python

5 `rc = myclient.group_leave(grp:string, directives:list)`

Python

6 **IN** `grp`

7 Python string identifier for the group (str)

8 **IN** `directives`

9 Optional list of Python [info](#) describing request (list)

10 Returns:

- 11 • `rc` - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

12 See [PMIx_Group_leave](#) for description of all relevant attributes and behaviors.

13 A.4.30 Client.group_destruct

14 Summary

15 Destruct a PMIx Group.

16 *PMIx v4.0* Format

Python

17 `rc = myclient.group_destruct(grp:string, directives:list)`

Python

18 **IN** `grp`

19 Python string identifier for the group (str)

20 **IN** `directives`

21 Optional list of Python [info](#) describing request (list)

22 Returns:

- 23 • `rc` - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

24 See [PMIx_Group_destruct](#) for description of all relevant attributes and behaviors.

25 A.4.31 Client.register_event_handler

26 Summary

27 Register an event handler to report events.

1

Format

Python

2

```
rc, id = myclient.register_event_handler(codes:list,  
3 directives:list, cbfunc)
```

3

Python

4

IN codes

List of Python integer status codes that should be reported to this handler (l1ist)

5

IN directives

Optional list of Python [info](#) describing request (list)

6

7

IN cbfunc

Python [evhandler](#) to be called when event is received (func)

8

9

Returns:

10

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *id* - PMIx reference identifier for handler (integer)

11

12

See [PMIx_Register_event_handler](#) for description of all relevant attributes and behaviors.

13

14 A.4.32 Client.deregister_event_handler

15

Summary

Deregister an event handler.

16

17 *PMIx v4.0*

Format

Python

18

```
myclient.deregister_event_handler(id:integer)
```

Python

19

IN id

PMIx reference identifier for handler (integer)

20

21

Returns: None

22

See [PMIx_Deregister_event_handler](#) for description of all relevant attributes and behaviors.

23

24 A.4.33 Client.notify_event

25

Summary

Report an event for notification via any registered handler.

26

Format

Python

```
rc = myclient.notify_event(status:integer, source:dict,  
                           range:integer, directives:list)
```

Python

IN status

PMIx status code indicating the event being reported (integer)

IN source

Python [proc](#) of the process that generated the event (dict)

IN range

Python [range](#) in which the event is to be reported (integer)

IN directives

Optional list of Python [info](#) dictionaries describing the event (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_Notify_event](#) for description of all relevant attributes and behaviors.

A.4.34 Client.fabric_register

Summary

Register for access to fabric-related information, including communication cost matrix.

Format

Python

```
rc,idx,fabricinfo = myclient.fabric_register(directives:list)
```

Python

IN directives

Optional list of Python [info](#) containing directives (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *idx* - Index of the registered fabric (integer)
- *fabricinfo* - List of Python [info](#) containing fabric info (list)

See [PMIx_Fabric_register](#) for details.

A.4.35 Client.fabric_update

Summary

Update fabric-related information, including communication cost matrix.

Format

Python

```
rc, fabricinfo = myclient.fabric_update(idx:integer)
```

Python

IN `idx`

Index of the registered fabric (list)

Returns:

- `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- `fabricinfo` - List of Python **info** containing updated fabric info (list)

See **PMIx_Fabric_update** for details.

A.4.36 Client.fabric_deregister

Summary

Deregister fabric.

Format

Python

PMIx v4.0

```
rc = myclient.fabric_deregister(idx:integer)
```

Python

IN `idx`

Index of the registered fabric (list)

Returns:

- `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See **PMIx_Fabric_deregister** for details.

A.4.37 Client.load_topology

Summary

Load the local hardware topology into the PMIx library.

Format

Python

PMIx v4.0

```
rc = myclient.load_topology()
```

Python

Returns:

- `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See **PMIx_Load_topology** for details - note that the topology loaded into the PMIx library may be utilized by PMIx and other libraries, but is not directly accessible by Python.

1 A.4.38 Client.get_relative_locality

2 Summary

3 Get the relative locality of two local processes.

4 *PMIx v4.0* Format

Python

5 `rc, locality = myclient.get_relative_locality(loc1:str, loc2:str)`

Python

6 **IN** `loc1`

7 Locality string of a process (str)

8 **IN** `loc2`

9 Locality string of a process (str)

10 Returns:

- 11 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 12 • `locality` - **locality** list containing the relative locality of the two processes (list)

13 See **PMIx_Get_relative_locality** for details.

14 A.4.39 Client.get_cpuset

15 Summary

16 Get the PU binding bitmap of the current process.

17 *PMIx v4.0* Format

Python

18 `rc, cpuset = myclient.get_cpuset(ref:integer)`

Python

19 **IN** `ref`

20 **bindenv** binding envelope to be used (integer)

21 Returns:

- 22 • `rc` - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- 23 • `cpuset` - **cpuset** containing the source and bitmap of the cpuset (dict)

24 See **PMIx_Get_cpuset** for details.

25 A.4.40 Client.parse_cpuset_string

26 Summary

27 Parse the PU binding bitmap from its string representation.

Format

Python

```
rc, cpuset = myclient.parse_cpuset_string(cpuset:string)
```

Python

IN `cpuset`

String returned by `PMIxServer.generate_cpuset_string` (string)

Returns:

- `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)
- `cpuset` - `cpuset` containing the source and bitmap of the cpuset (dict)

See `PMIx_Parse_cpuset_string` for details.

A.4.41 Client.compute_distances

Summary

Compute distances from specified process location to local devices.

Format

Python

```
rc, distances = myclient.compute_distances(cpuset:dict, info:list)
```

Python

IN `cpuset`

`cpuset` describing the location of the process (dict)

IN `info`

List of `info` dictionaries describing the devices whose distance is to be computed (list)

Returns:

- `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)
- `distances` - List of `devdist` structures containing the distances from the caller to the specified devices (list)

See `PMIx_Compute_distances` for details. Note that distances can only be computed against the local topology.

A.4.42 Client.error_string

Summary

Pretty-print string representation of `pmix_status_t`.

1 **Format** Python

2 `rep = myclient.error_string(status:integer)`
3 Python

4 **IN** `status`
5 PMIx status code (integer)

6 Returns:
7 • `rep` - String representation of the provided status code (str)
8 See [PMIx_Error_string](#) for further details.

8 **A.4.43 Client.proc_state_string**

9 **Summary**
10 Pretty-print string representation of [pmix_proc_state_t](#).

11 *PMIx v4.0* **Format** Python

12 `rep = myclient.proc_state_string(state:integer)`
13 Python

14 **IN** `state`
15 PMIx process state code (integer)

16 Returns:
17 • `rep` - String representation of the provided process state (str)
18 See [PMIx_Proc_state_string](#) for further details.

18 **A.4.44 Client.scope_string**

19 **Summary**
20 Pretty-print string representation of [pmix_scope_t](#).

21 *PMIx v4.0* **Format** Python

22 `rep = myclient.scope_string(scope:integer)`
23 Python

24 **IN** `scope`
25 PMIx scope value (integer)

26 Returns:
27 • `rep` - String representation of the provided scope (str)
28 See [PMIx_Scope_string](#) for further details

1 A.4.45 Client.persistence_string

2 Summary

3 Pretty-print string representation of [pmix_persistence_t](#).

4 *PMIx v4.0* Format

Python

5 `rep = myclient.persistence_string(persistence:integer)`

Python

6 IN persistence

7 PMIx persistence value (integer)

8 Returns:

- 9 • *rep* - String representation of the provided persistence (str)

10 See [PMIx_Persistence_string](#) for further details.

11 A.4.46 Client.data_range_string

12 Summary

13 Pretty-print string representation of [pmix_data_range_t](#).

14 *PMIx v4.0* Format

Python

15 `rep = myclient.data_range_string(range:integer)`

Python

16 IN range

17 PMIx data range value (integer)

18 Returns:

- 19 • *rep* - String representation of the provided data range (str)

20 See [PMIx_Data_range_string](#) for further details.

21 A.4.47 Client.info_directives_string

22 Summary

23 Pretty-print string representation of [pmix_info_directives_t](#).

1 **Format** Python

2 `rep = myclient.info_directives_string(directives:bitarray)`
3 Python

4 **IN directives**
5 PMIx **info directives** value (bitarray)

6 Returns:
7 • *rep* - String representation of the provided info directives (str)
8 See [PMIx_Info_directives_string](#) for further details.

8 **A.4.48 Client.data_type_string**

9 **Summary**
10 Pretty-print string representation of [pmix_data_type_t](#).

11 *PMIx v4.0* **Format** Python

12 `rep = myclient.data_type_string(dtype:integer)`
13 Python

14 **IN dtype**
15 PMIx datatype value (integer)

16 Returns:
17 • *rep* - String representation of the provided datatype (str)
18 See [PMIx_Data_type_string](#) for further details.

18 **A.4.49 Client.alloc_directive_string**

19 **Summary**
20 Pretty-print string representation of [pmix_alloc_directive_t](#).

21 *PMIx v4.0* **Format** Python

22 `rep = myclient.alloc_directive_string(adir:integer)`
23 Python

24 **IN adir**
25 PMIx allocation directive value (integer)

26 Returns:
27 • *rep* - String representation of the provided allocation directive (str)
28 See [PMIx_Alloc_directive_string](#) for further details.

1 A.4.50 Client.iof_channel_string

2 Summary

3 Pretty-print string representation of [pmix_iof_channel_t](#).

4 Format

PMIx v4.0

Python

```
5 rep = myclient.iof_channel_string(channel:bitarray)
```

Python

6 **IN** channel

7 PMIx IOF [channel](#) value (bitarray)

8 Returns:

- 9 • *rep* - String representation of the provided IOF channel (str)

10 See [PMIx_IOF_channel_string](#) for further details.

11 A.4.51 Client.job_state_string

12 Summary

13 Pretty-print string representation of [pmix_job_state_t](#).

14 Format

PMIx v4.0

Python

```
15 rep = myclient.job_state_string(state:integer)
```

Python

16 **IN** state

17 PMIx job state value (integer)

18 Returns:

- 19 • *rep* - String representation of the provided job state (str)

20 See [PMIx_Job_state_string](#) for further details.

21 A.4.52 Client.get_attribute_string

22 Summary

23 Pretty-print string representation of a PMIx attribute.

Format

Python

```
rep = myclient.get_attribute_string(attribute: str)
```

Python

IN attribute

PMIx attribute name (string)

Returns:

- *rep* - String representation of the provided attribute (str)
- See [PMIx_Get_attribute_string](#) for further details.

A.4.53 Client.get_attribute_name

Summary

Pretty-print name of a PMIx attribute corresponding to the provided string.

Format

Python

PMIx v4.0

```
rep = myclient.get_attribute_name(attribute: str)
```

Python

IN attributestring

Attribute string (string)

Returns:

- *rep* - Attribute name corresponding to the provided string (str)
- See [PMIx_Get_attribute_name](#) for further details.

A.4.54 Client.link_state_string

Summary

Pretty-print string representation of [pmix_link_state_t](#).

Format

Python

PMIx v4.0

```
rep = myclient.link_state_string(state: integer)
```

Python

IN state

PMIx link state value (integer)

Returns:

- *rep* - String representation of the provided link state (str)
- See [PMIx_Link_state_string](#) for further details.

1 A.4.55 Client.device_type_string

2 Summary

3 Pretty-print string representation of [pmix_device_type_t](#).

4 *PMIx v4.0* Format

Python

5 `rep = myclient.device_type_string(type:bitarray)`

Python

6 IN type

7 PMIx device type value (bitarray)

8 Returns:

- 9 • *rep* - String representation of the provided device type (str)

10 See [PMIx_Device_type_string](#) for further details.

11 A.4.56 Client.progress

12 Summary

13 Progress the PMIx library.

14 *PMIx v4.0* Format

Python

15 `myclient.progress()`

Python

16 See [PMIx_Progress](#) for further details.

17 A.5 PMIxServer

18 The server Python class inherits the Python "client" class as its parent. Thus, it includes all client
19 functions in addition to the ones defined in this section.

20 A.5.1 Server.init

21 Summary

22 Initialize the PMIx server library after obtaining a new PMIxServer object.

1
2
3
4
5
6
7
8
9
10

Format

Python

```
rc = myserver.init(directives:list, map:dict)
```

Python

IN directives

List of Python [info](#) dictionaries (list)

IN map

Python dictionary key-function pairs that map [server module](#) callback functions to provided implementations (see [pmix_server_module_t](#)) (dict)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- See [PMIx_server_init](#) for description of all relevant attributes and behaviors.

11 **A.5.2 Server.finalize**

12 **Summary**

13 Finalize the PMIx server library.

14 *PMIx v4.0*

Format

Python

```
rc = myserver.finalize()
```

Python

16 Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- 17 See [PMIx_server_finalize](#) for details.

19 **A.5.3 Server.generate_regex**

20 **Summary**

21 Generate a regular expression representation of the input strings.

Format

Python

```
rc, regex = myserver.generate_regex(input:list)
```

Python

IN input

List of Python strings (e.g., node names) (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *regex* - Python **bytearray** containing regular expression representation of the input list (**bytearray**)

See [PMIx_generate_regex](#) for details.

A.5.4 Server.generate_ppn

Summary

Generate a regular expression representation of the input strings.

Format

Python

```
rc, regex = myserver.generate_ppn(input:list)
```

Python

IN input

List of Python strings, each string consisting of a comma-delimited list of ranks on each node, with the strings being in the same order as the node names provided to "generate_regex" (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *regex* - Python **bytearray** containing regular expression representation of the input list (**bytearray**)

See [PMIx_generate_ppn](#) for details.

A.5.5 Server.generate_locality_string

Summary

Generate a PMIx locality string from a given cpuset.

Format

Python

```
rc, locality = myserver.generate_locality_string(cpuset:dict)
```

Python

IN cset

cpuset containing the bitmap of assigned PUs (dict)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *locality* - String representation of the PMIx locality corresponding to the input bitmap (string)

See [PMIx_server_generate_locality_string](#) for details.

A.5.6 Server.generate_cpuset_string

Summary

Generate a PMIx string representation of the provided cpuset.

Format

Python

```
rc, cpustr = myserver.generate_cpuset_string(cpuset:dict)
```

Python

IN cset

cpuset containing the bitmap of assigned PUs (dict)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *cpustr* - String representation of the input bitmap (string)

See [PMIx_server_generate_cpuset_string](#) for details.

A.5.7 Server.register_namespace

Summary

Setup the data about a particular namespace.

Format

Python

```
rc = myserver.register_namespace(namespace: str,  
                                nlocalprocs: integer,  
                                directives: list)
```

Python

IN namespace

Python string containing the namespace (str)

IN nlocalprocs

Number of local processes (integer)

IN directives

List of Python [info](#) dictionaries (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_server_register_namespace](#) for description of all relevant attributes and behaviors.

A.5.8 Server.deregister_namespace

Summary

Deregister a namespace.

Format

Python

```
myserver.deregister_namespace(namespace: str)
```

Python

IN namespace

Python string containing the namespace (str)

Returns: None

See [PMIx_server_deregister_namespace](#) for details.

A.5.9 Server.register_resources

Summary

Register non-namespace related information with the local PMIx library

1 **Format** Python

2 `myserver.register_resources(directives:list)` Python

3 **IN directives**
4 List of Python **info** dictionaries (list)

5 Returns: None

6 See [PMIx_server_register_resources](#) for details.

7 **A.5.10 Server.deregister_resources**

8 **Summary**

9 Remove non-namespace related information from the local PMIx library

10 *PMIx v4.0* **Format** Python

11 `myserver.deregister_resources(directives:list)` Python

12 **IN directives**
13 List of Python **info** dictionaries (list)

14 Returns: None

15 See [PMIx_server_deregister_resources](#) for details.

16 **A.5.11 Server.register_client**

17 **Summary**

18 Register a client process with the PMIx server library.

19 *PMIx v4.0* **Format** Python

20 `rc = myserver.register_client(proc:dict, uid:integer, gid:integer)` Python

21 **IN proc**
22 Python **proc** dictionary identifying the client process (dict)

23 **IN uid**
24 Linux uid value for user executing client process (integer)

25 **IN gid**
26 Linux gid value for user executing client process (integer)

27 Returns:

- 28 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

29 See [PMIx_server_register_client](#) for details.

1 A.5.12 Server.deregister_client

2 Summary

3 Deregister a client process and purge all data relating to it.

4 *PMIx v4.0* Format

Python

5 `myserver.deregister_client(proc:dict)`

Python

6 **IN** `proc`

7 Python `proc` dictionary identifying the client process (dict)

8 Returns: None

9 See [PMIx_server_deregister_client](#) for details.

10 A.5.13 Server.setup_fork

11 Summary

12 Setup the environment of a child process that is to be forked by the host.

13 *PMIx v4.0* Format

Python

14 `rc = myserver.setup_fork(proc:dict, environ:dict)`

Python

15 **IN** `proc`

16 Python `proc` dictionary identifying the client process (dict)

17 **INOUT** `environ`

18 Python dictionary containing the environment to be passed to the client (dict)

19 Returns:

- 20 • `rc` - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

21 See [PMIx_server_setup_fork](#) for details.

22 A.5.14 Server.dmodex_request

23 Summary

24 Function by which the host server can request modex data from the local PMIx server.

Format

Python

```
rc, data = myserver.dmodex_request(proc:dict)
```

Python

IN `proc`

Python `proc` dictionary identifying the process whose data is requested (dict)

Returns:

- `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)
- `data` - Python `byteobject` containing the returned data (dict)

See `PMIx_server_dmodex_request` for details.

A.5.15 `Server.setup_application`

Summary

Function by which the resource manager can request application-specific setup data prior to launch of a *job*.

Format

Python

```
rc, info = myserver.setup_application(namespace:str, directives:list)
```

Python

IN `namespace`

Namespace whose setup information is being requested (str)

IN `directives`

Python list of `info` directives

Returns:

- `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)
- `info` - Python list of `info` dictionaries containing the returned data (list)

See `PMIx_server_setup_application` for details.

A.5.16 `Server.register_attributes`

Summary

Register host environment attribute support for a function.

1
2
3
4
5
6
7
8
9

Format

Python

```
rc = myserver.register_attributes(function:str, attrs:list)
```

Python

IN function

Name of the function (str)

IN attrs

Python list of [regattr](#) describing the supported attributes

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_Register_attributes](#) for details.

10 A.5.17 Server.setup_local_support

11 Summary

12 Function by which the local PMIx server can perform any application-specific operations prior to
13 spawning local clients of a given application.

14 *PMIx v4.0* **Format**

Python

```
rc = myserver.setup_local_support(nspace:str, info:list)
```

Python

16 **IN nspace**

17 Namespace whose setup information is being requested (str)

18 **IN info**

19 Python list of [info](#) containing the setup data (list)

20 Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

22 See [PMIx_server_setup_local_support](#) for details.

23 A.5.18 Server.iof_deliver

24 Summary

25 Function by which the host environment can pass forwarded IO to the PMIx server library for
26 distribution to its clients.

Format

Python

```
rc = myserver.iof_deliver(source:dict, channel:integer,  
                          data:dict, directives:list)
```

Python

IN source

Python **proc** dictionary identifying the process who generated the data (dict)

IN channel

Python **channel** bitmask identifying IO channel of the provided data (integer)

IN data

Python **byteobject** containing the data (dict)

IN directives

Python list of **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_server_IOF_deliver](#) for details.

A.5.19 Server.collect_inventory

Summary

Collect inventory of resources on a node.

Format

Python

```
rc, info = myserver.collect_inventory(directives:list)
```

Python

IN directives

Optional Python list of **info** containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *info* - Python list of **info** containing the returned data (list)

See [PMIx_server_collect_inventory](#) for details.

A.5.20 Server.deliver_inventory

Summary

Pass collected inventory to the PMIx server library for storage.

Format

Python

```
rc = myserver.deliver_inventory(info:list, directives:list)
```

Python

IN info

- Python list of **info** dictionaries containing the inventory data (list)

IN directives

Python list of **info** dictionaries containing directives (list)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_server_deliver_inventory](#) for details.

A.5.21 Server.define_process_set

Summary

Add members to a PMIx process set.

Format

Python

```
rc = myserver.define_process_set(members:list, name:str)
```

Python

IN members

- List of Python **proc** dictionaries identifying the processes to be added to the process set (list)

IN name

- Name of the process set (str)

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_server_define_process_set](#) for details.

A.5.22 Server.delete_process_set

Summary

Delete a PMIx process set.

1 **Format** Python

```
2 rc = myserver.delete_process_set(name:str)
```

Python

3 **IN** name
4 - Name of the process set (str)

5 Returns:
6 • rc - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
7 See [PMIx_server_delete_process_set](#) for details.

8 **A.5.23 Server.register_resources**

9 **Summary**
10 Register non-namespace related information with the local PMIx server library.

11 *PMIx v4.0* **Format** Python

```
12 rc = myserver.register_resources(info:list)
```

Python

13 **IN** info
14 - List of Python **info** dictionaries list)

15 Returns:
16 • rc - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
17 See [PMIx_server_register_resources](#) for details.

18 **A.5.24 Server.deregister_resources**

19 **Summary**
20 Deregister non-namespace related information with the local PMIx server library.

21 *PMIx v4.0* **Format** Python

```
22 rc = myserver.deregister_resources(info:list)
```

Python

23 **IN** info
24 - List of Python **info** dictionaries list)

25 Returns:
26 • rc - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
27 See [PMIx_server_deregister_resources](#) for details.

1 A.6 PMIxTool

2 The tool Python class inherits the Python "server" class as its parent. Thus, it includes all client and
3 server functions in addition to the ones defined in this section.

4 A.6.1 Tool.init

5 Summary

6 Initialize the PMIx tool library after obtaining a new PMIxTool object.

7 *PMIx v4.0* Format

Python

8 `rc,proc = mytool.init(info:list)`

Python

9 IN info

10 List of Python `info` directives (list)

11 Returns:

- 12 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)
- 13 • `proc` - a Python `proc` (dict)

14 See `PMIx_tool_init` for description of all relevant attributes and behaviors.

15 A.6.2 Tool.finalize

16 Summary

17 Finalize the PMIx tool library, closing the connection to the server.

18 *PMIx v4.0* Format

Python

19 `rc = mytool.finalize()`

Python

20 Returns:

- 21 • `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

22 See `PMIx_tool_finalize` for description of all relevant attributes and behaviors.

23 A.6.3 Tool.disconnect

24 Summary

25 Disconnect the PMIx tool from the specified server connection while leaving the tool library
26 initialized.

1 **Format** Python

2 `rc = mytool.disconnect(server:dict)` Python

3 **IN server**
4 Process identifier of server from which the tool is to be disconnected (**proc**)

5 Returns:
6 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
7 See **PMIx_tool_disconnect** for details.

8 A.6.4 Tool.attach_to_server

9 **Summary**
10 Establish a connection to a PMIx server.

11 *PMIx v4.0* **Format** Python

12 `rc,proc,server = mytool.connect_to_server(info:list)` Python

13 **IN info**
14 List of Python **info** dictionaries (list)

15 Returns:
16 • *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
17 • *proc* - a Python **proc** containing the tool's identifier (dict)
18 • *server* - a Python **proc** containing the identifier of the server to which the tool attached (dict)
19 See **PMIx_tool_attach_to_server** for details.

20 A.6.5 Tool.get_servers

21 **Summary**
22 Get a list containing the **proc** process identifiers of all servers to which the tool is currently
23 connected.

1
2
3
4
5
6
7

Format

Python

```
rc, servers = mytool.get_servers()
```

Python

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)
- *servers* - a list of Python **proc** containing the identifiers of the servers to which the tool is currently attached (dict)

See **PMIx_tool_get_servers** for details.

8 **A.6.6 Tool.set_server**

9

Summary

Designate a server as the tool's *primary* server.

11 *PMIx v4.0*

Format

Python

```
rc = mytool.set_server(proc:dict, info:list)
```

Python

13
14
15
16

IN proc

Python **proc** containing the identifier of the servers to which the tool is to attach (list)

IN info

List of Python **info** dictionaries (list)

17
18
19

Returns:

- *rc* - **PMIX_SUCCESS** or a negative value corresponding to a PMIx error constant (integer)

See **PMIx_tool_set_server** for details.

20 **A.6.7 Tool.iof_pull**

21
22

Summary

Register to receive output forwarded from a remote process.

Format

Python

```
rc, id = mytool.iof_pull(sources:list, channel:integer,  
                        directives:list, cbfunc)
```

Python

IN sources

List of Python [proc](#) dictionaries of processes whose IO is being requested (list)

IN channel

Python [channel](#) bitmask identifying IO channels to be forwarded (integer)

IN directives

List of Python [info](#) dictionaries describing request (list)

IN cbfunc

Python [iofcbfunc](#) to receive IO payloads (func)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)
- *id* - PMIx reference identifier for request (integer)

See [PMIx_IOF_pull](#) for description of all relevant attributes and behaviors.

A.6.8 Tool.iof_deregister

Summary

Deregister from output forwarded from a remote process.

Format

Python

```
rc = mytool.iof_deregister(id:integer, directives:list)
```

Python

IN id

PMIx reference identifier returned by pull request (list)

IN directives

List of Python [info](#) dictionaries describing request (list)

Returns:

- *rc* - [PMIX_SUCCESS](#) or a negative value corresponding to a PMIx error constant (integer)

See [PMIx_IOF_deregister](#) for description of all relevant attributes and behaviors.

A.6.9 Tool.iof_push

Summary

Push data collected locally (typically from stdin) to stdin of target recipients.

Format

Python

```
rc = mytool.iof_push(targets:list, data:dict, directives:list)
```

Python

IN sources

List of Python `proc` of target processes (list)

IN data

Python `byteobject` containing data to be delivered (dict)

IN directives

Optional list of Python `info` describing request (list)

Returns:

- `rc` - `PMIX_SUCCESS` or a negative value corresponding to a PMIx error constant (integer)

See `PMIX_IOF_push` for description of all relevant attributes and behaviors.

A.7 Example Usage

The following examples are provided to illustrate the use of the Python bindings.

A.7.1 Python Client

The following example contains a client program that illustrates a fairly common usage pattern.

The program instantiates and initializes the `PMIxClient` class, posts some data that is to be shared across all processes in the job, executes a “fence” that circulates the data, and then retrieves a value posted by one of its peers. Note that the example has been formatted to fit the document layout.

Python

```
from pmix import *

def main():
    # Instantiate a client object
    myclient = PMIxClient()
    print("Testing PMIx ", myclient.get_version())

    # Initialize the PMIx client library, declaring the programming model
    # as "TEST" and the library name as "PMIX", just for the example
    info = ['key':PMIX_PROGRAMMING_MODEL,
            'value':'TEST', 'val_type':PMIX_STRING,
            'key':PMIX_MODEL_LIBRARY_NAME,
            'value':'PMIX', 'val_type':PMIX_STRING]
    rc,myname = myclient.init(info)
```

```

1     if PMIX_SUCCESS != rc:
2         print("FAILED TO INIT WITH ERROR", myclient.error_string(rc))
3         exit(1)
4
5     # try posting a value
6     rc = myclient.put(PMIX_GLOBAL, "mykey",
7                       'value':1, 'val_type':PMIX_INT32)
8     if PMIX_SUCCESS != rc:
9         print("PMIx_Put FAILED WITH ERROR", myclient.error_string(rc))
10        # cleanly finalize
11        myclient.finalize()
12        exit(1)
13
14    # commit it
15    rc = myclient.commit()
16    if PMIX_SUCCESS != rc:
17        print("PMIx_Commit FAILED WITH ERROR",
18              myclient.error_string(rc))
19        # cleanly finalize
20        myclient.finalize()
21        exit(1)
22
23    # execute fence across all processes in my job
24    procs = []
25    info = []
26    rc = myclient.fence(procs, info)
27    if PMIX_SUCCESS != rc:
28        print("PMIx_Fence FAILED WITH ERROR", myclient.error_string(rc))
29        # cleanly finalize
30        myclient.finalize()
31        exit(1)
32
33    # Get a value from a peer
34    if 0 != myname['rank']:
35        info = []
36        rc, get_val = myclient.get('nspace':"testnspace", 'rank': 0,
37                                  "mykey", info)
38    if PMIX_SUCCESS != rc:
39        print("PMIx_Commit FAILED WITH ERROR",
40              myclient.error_string(rc))
41        # cleanly finalize
42        myclient.finalize()
43        exit(1)

```

```

1         print("Get value returned: ", get_val)
2
3     # test a fence that should return not_supported because
4     # we pass a required attribute that the server is known
5     # not to support
6     procs = []
7     info = ['key': 'ARBIT', 'flags': PMIX_INFO_REQD,
8             'value':10, 'val_type':PMIX_INT]
9     rc = myclient.fence(procs, info)
10    if PMIX_SUCCESS == rc:
11        print("PMIx_Fence SUCCEEDED BUT SHOULD HAVE FAILED")
12        # cleanly finalize
13        myclient.finalize()
14        exit(1)
15
16    # Publish something
17    info = ['key': 'ARBITRARY', 'value':10, 'val_type':PMIX_INT]
18    rc = myclient.publish(info)
19    if PMIX_SUCCESS != rc:
20        print("PMIx_Publish FAILED WITH ERROR",
21              myclient.error_string(rc))
22        # cleanly finalize
23        myclient.finalize()
24        exit(1)
25
26    # finalize
27    info = []
28    myclient.finalize(info)
29    print("Client finalize complete")
30
31    # Python main program entry point
32    if __name__ == '__main__':
33        main()

```

Python

34 A.7.2 Python Server

35 The following example contains a minimum-level server host program that instantiates and
36 initializes the `PMIxServer` class. The program illustrates passing several server module functions to
37 the bindings and includes code to setup and spawn a simple client application, waiting until the
38 spawned client terminates before finalizing and exiting itself. Note that the example has been
39 formatted to fit the document layout.

```
1 from pmix import *
2 import signal, time
3 import os
4 import select
5 import subprocess
6
7 def clientconnected(proc:tuple is not None):
8     print("CLIENT CONNECTED", proc)
9     return PMIX_OPERATION_SUCCEEDED
10
11 def clientfinalized(proc:tuple is not None):
12     print("CLIENT FINALIZED", proc)
13     return PMIX_OPERATION_SUCCEEDED
14
15 def clientfence(procs:list, directives:list, data:bytearray):
16     # check directives
17     if directives is not None:
18         for d in directives:
19             # these are each an info dict
20             if "pmix" not in d['key']:
21                 # we do not support such directives - see if
22                 # it is required
23                 try:
24                     if d['flags'] & PMIX_INFO_REQD:
25                         # return an error
26                         return PMIX_ERR_NOT_SUPPORTED
27                 except:
28                     #it can be ignored
29                     pass
30     return PMIX_OPERATION_SUCCEEDED
31
32 def main():
33     try:
34         myserver = PMIXServer()
35     except:
36         print("FAILED TO CREATE SERVER")
37         exit(1)
38     print("Testing server version ", myserver.get_version())
39
40     args = ['key':PMIX_SERVER_SCHEDULER,
41            'value':'T', 'val_type':PMIX_BOOL]
42     map = 'clientconnected': clientconnected,
```

```

1         'clientfinalized': clientfinalized,
2         'fencenb': clientfence
3 my_result = myserver.init(args, map)
4
5 # get our environment as a base
6 env = os.environ.copy()
7
8 # register an nspace for the client app
9 (rc, regex) = myserver.generate_regex("test000,test001,test002")
10 (rc, ppn) = myserver.generate_ppn("0")
11 kvals = ['key':PMIX_NODE_MAP,
12         'value':regex, 'val_type':PMIX_STRING,
13         'key':PMIX_PROC_MAP,
14         'value':ppn, 'val_type':PMIX_STRING,
15         'key':PMIX_UNIV_SIZE,
16         'value':1, 'val_type':PMIX_UINT32,
17         'key':PMIX_JOB_SIZE,
18         'value':1, 'val_type':PMIX_UINT32]
19 rc = foo.register_nspace("testnspace", 1, kvals)
20 print("RegNspace ", rc)
21
22 # register a client
23 uid = os.getuid()
24 gid = os.getgid()
25 rc = myserver.register_client('nspace':"testnspace", 'rank':0,
26                             uid, gid)
27 print("RegClient ", rc)
28 # setup the fork
29 rc = myserver.setup_fork('nspace':"testnspace", 'rank':0, env)
30 print("SetupFrk", rc)
31
32 # setup the client argv
33 args = ["/client.py"]
34 # open a subprocess with stdout and stderr
35 # as distinct pipes so we can capture their
36 # output as the process runs
37 p = subprocess.Popen(args, env=env,
38                     stdout=subprocess.PIPE, stderr=subprocess.PIPE)
39 # define storage to catch the output
40 stdout = []
41 stderr = []
42 # loop until the pipes close
43 while True:

```

```

1     reads = [p.stdout_FILENO(), p.stderr_FILENO()]
2     ret = select.select(reads, [], [])
3
4     stdout_done = True
5     stderr_done = True
6
7     for fd in ret[0]:
8         # if the data
9         if fd == p.stdout_FILENO():
10            read = p.stdout.readline()
11            if read:
12                read = read.decode('utf-8').rstrip()
13                print('stdout: ' + read)
14                stdout_done = False
15            elif fd == p.stderr_FILENO():
16                read = p.stderr.readline()
17                if read:
18                    read = read.decode('utf-8').rstrip()
19                    print('stderr: ' + read)
20                    stderr_done = False
21
22            if stdout_done and stderr_done:
23                break
24            print("FINALIZING")
25            myserver.finalize()
26
27
28 if __name__ == '__main__':
29     main()

```

Python

APPENDIX B

Use-Cases

1 The PMIx standard provides many generic interfaces that can be composed into higher-level use
2 cases in a variety of ways. While the specific interfaces and attributes are standardized, the use
3 cases themselves are not (and should not) be standardized. Common use cases are included here as
4 examples of how PMIx's generic interfaces *might* be composed together for a higher-level purpose.
5 The use cases are intended for both PMIx interface users and library implementors. Whereby a
6 better understanding of the general usage model within the community can help users picking up
7 PMIx for the first and help implementors optimize their implementation for the common cases.

8 Each use case is structured to provide background information about the high-level use case as well
9 as specific details about how the PMIx interfaces are used within the use case. Some use cases even
10 provide code snippets. These code snippets are apart of larger code examples located within the
11 standard's source code repository, and each complete code example is fully compilable and
12 runnable. The related interfaces and attributes collected at the bottom of each use case are mainly
13 for convenience and link to the full standardized definitions.

14 **B.1 Business Card Exchange for Process-to-** 15 **Process Wire-up**

16 **B.1.1 Use Case Summary**

17 Multi-process communication libraries, such as MPI, need to establish communication channels
18 between a set of those processes. In this scenario, each process needs to share connectivity
19 information (a.k.a. Business Cards) with all other processes before communication channels can be
20 established. This connectivity information may take the form of one or more unique strings that
21 allow a different process to establish a communication channel with the originator. The runtime
22 environment must provide a mechanism for the efficient exchange of this connectivity information.
23 Additional information about the current state of the job (e.g., number of processes globally and
24 locally) and of how the process was started (e.g., process binding) is also helpful.

25 Note: The Instant-On wire-up mechanism is a separate, related use case.

1 B.1.2 Use Case Details

2 Each process provides their business card to PMIx via one or more **PMIx_Put** operations to store
3 the tuple of {**UID**, **key**, **value**}. The **UID** is the unique name for this process in the PMIx
4 universe (i.e., **namespace** and **rank**). The **key** is a unique key that other processes can reference
5 generically (note that since the **UID** is also associated with the **key** there is no need to make the
6 **key** uniquely named per process). The **value** is the string representation of the connectivity
7 information.

8 Some business card information is meant for remote processes (e.g., TCP or InfiniBand addresses)
9 while others are meant only for local processes (e.g., shared memory information). As such a
10 **scope** should be associated with the **PMIx_Put** operation to differentiate this intention.

11 The **PMIx_Put** operations may be cached local to the process. Once all **PMIx_Put** operations
12 have been called each process should call **PMIx_Commit** to push those values to the local PMIx
13 server. Note that in a multi-library configuration each library may **PMIx_Put** then
14 **PMIx_Commit** values - so there may be multiple **PMIx_Commit** calls before a Business Card
15 Exchange is activated.

16 After calling **PMIx_Commit** a process can activate the Business Card Exchange collective
17 operation by calling **PMIx_Fence**. The **PMIx_Fence** operation is collective over the set of
18 processes specified in the argument set. That allows for the collective to span a subset of a
19 namespace or multiple namespaces. After the completion of the **PMIx_Fence** operation, the data
20 stored by other processes via **PMIx_Put** is available to the local process through a call to
21 **PMIx_Get** which returns the key/value pairs necessary to establish the connection(s) with the
22 other processes.

23 The **PMIx_Fence** operation has a "Synchronize Only" mode that works as a barrier operation.
24 This is helpful if the communication library requires a synchronization before leaving initialization
25 or starting finalization, for example.

26 The **PMIx_Fence** operation has a "Sparse" mode in addition to a "Full" mode for the data
27 exchange. The "Full" mode will fully exchange all Business Card information with all other
28 processes. This is helpful for tightly communicating applications. The "Sparse" mode will
29 dynamically pull the connectivity information on-demand from inside of **PMIx_Get** (if it is not
30 already available locally). This is helpful for sparsely communicating applications. Since which
31 mode is best for an application cannot be inferred by the PMIx library the caller must specify which
32 mode works best for their application. The **PMIx_Fence** operation has an option for the end user
33 to specify which mode they desire for this operation.

34 Additional information about the current state of the job (e.g., number of processes globally and
35 locally) and of how the process was started (e.g., process binding) is also helpful. This "job level"
36 information is available immediately after **PMIx_Init** without the need for any explicit
37 synchronization.

38 The number of processes globally in the namespace and this process's rank within that namespace
39 is important to know before establishing the Business Card information to best allocate resources.

1 The number of processes local to the node and this process's local rank is important to know before
2 establishing the Business Card information to help the caller determine the scope of the put
3 operation. For example, to designate a leader to set up a shared memory segment of the proper size
4 before putting that information into the locally scoped Business Card information.

5 The number of processes local to a remote node is also helpful to know before establishing the
6 Business Card information. This information is useful to pre-establish local resources before that
7 remote node starts to initiate a connection or to determine the number of connections that need to
8 be advertised in the Business Card when it is sent out.

9 Note that some of the job level information may change over the course of the job in a dynamic
10 application.

11 Related Interfaces

12 *PMIx v1.0* **PMIx_Put**

▼  ▼

13 `pmix_status_t`
14 `PMIx_Put (pmix_scope_t scope,`
15 `const pmix_key_t key,`
16 `pmix_value_t *val);`

▲  ▲

17 *PMIx v1.0* **PMIx_Get**

▼  ▼

18 `pmix_status_t`
19 `PMIx_Get (const pmix_proc_t *proc, const pmix_key_t key,`
20 `const pmix_info_t info[], size_t ninfo,`
21 `pmix_value_t **val);`

▲  ▲

22 *PMIx v1.0* **PMIx_Commit**

▼  ▼

23 `pmix_status_t PMIx_Commit (void);`

▲  ▲

24 *PMIx v1.0* **PMIx_Fence**

▼  ▼

25 `pmix_status_t`
26 `PMIx_Fence (const pmix_proc_t procs[], size_t nprocs,`
27 `const pmix_info_t info[], size_t ninfo);`

PMIx_Init

```
pmix_status_t
PMIx_Init (pmix_proc_t *proc,
           pmix_info_t info[], size_t ninfo)
```

Related Attributes

The following job level information is useful to have before establishing Business Card information:

PMIX_NODE_LIST "pmix.nlist" (char*)

Comma-delimited list of nodes currently hosting processes in the specified realm. Defaults to the *job* realm.

PMIX_NUM_NODES "pmix.num.nodes" (uint32_t)

Number of nodes currently hosting processes in the specified realm. Defaults to the *job* realm.

PMIX_NODEID "pmix.nodeid" (uint32_t)

Node identifier expressed as the node's index (beginning at zero) in an array of nodes within the active session. The value must be unique and directly correlate to the **PMIX_HOSTNAME** of the node - i.e., users can interchangeably reference the same location using either the **PMIX_HOSTNAME** or corresponding **PMIX_NODEID**.

PMIX_JOB_SIZE "pmix.job.size" (uint32_t)

Total number of processes in the specified job across all contained applications. Note that this value can be different from **PMIX_MAX_PROCS**. For example, users may choose to subdivide an allocation (running several jobs in parallel within it), and dynamic programming models may support adding and removing processes from a running *job* on-the-fly. In the latter case, PMIx events may be used to notify processes within the job that the job size has changed.

PMIX_PROC_MAP "pmix.pmap" (char*)

Regular expression describing processes on each node in the specified realm - see [17.2.3.2](#) for an explanation of its generation. Defaults to the *job* realm.

PMIX_LOCAL_PEERS "pmix.lpeers" (char*)

Comma-delimited list of ranks that are executing on the local node within the specified namespace – shortcut for **PMIx_Resolve_peers** for the local node.

PMIX_LOCAL_SIZE "pmix.local.size" (uint32_t)

Number of processes in the specified job or application on the caller's node. Defaults to job unless the **PMIX_APP_INFO** and the **PMIX_APPNUM** qualifiers are given.

1 For each process this information is also useful (note that any one process may want to access this
2 list of information about any other process in the system):

3 **PMIX_RANK** "pmix.rank" (pmix_rank_t)

4 Process rank within the job, starting from zero.

5 **PMIX_LOCAL_RANK** "pmix.lrank" (uint16_t)

6 Rank of the specified process on its node - refers to the numerical location (starting from
7 zero) of the process on its node when counting only those processes from the same job that
8 share the node, ordered by their overall rank within that job.

9 **PMIX_GLOBAL_RANK** "pmix.grank" (pmix_rank_t)

10 Rank of the specified process spanning across all jobs in this session, starting with zero.
11 Note that no ordering of the jobs is implied when computing this value. As jobs can start and
12 end at random times, this is defined as a continually growing number - i.e., it is not
13 dynamically adjusted as individual jobs and processes are started or terminated.

14 **PMIX_LOCALITY_STRING** "pmix.locstr" (char*)

15 String describing a process's bound location - referenced using the process's rank. The string
16 is prefixed by the implementation that created it (e.g., "hwloc") followed by a colon. The
17 remainder of the string represents the corresponding locality as expressed by the underlying
18 implementation. The entire string must be passed to **PMIx_Get_relative_locality**
19 for processing. Note that hosts are only required to provide locality strings for local client
20 processes - thus, a call to **PMIx_Get** for the locality string of a process that returns
21 **PMIX_ERR_NOT_FOUND** indicates that the process is not executing on the same node.

22 **PMIX_HOSTNAME** "pmix.hname" (char*)

23 Name of the host, as returned by the **gethostname** utility or its equivalent.

24 There are other keys that are helpful to have before a synchronization point. This is not meant to be
25 a comprehensive list.

26 B.2 Debugging

27 B.2.1 Terminology

28 B.2.1.1 Tools vs Debuggers

29 A *tool* is a process designed to monitor, record, analyze, or control the execution of another
30 process. Typically used for the purposes of profiling and debugging. A *first-party tool* runs within
31 the address space of the application process while a *third-party tool* run within its own process. A
32 *debugger* is a third-party tool that inspects and controls an application process's execution using
33 system-level debug APIs (e.g., **ptrace**).

1 B.2.1.2 Parallel Launching Methods

2 A *starter* program is a program responsible for launching a parallel runtime, such as MPI. PMIx
3 supports two primary methods for launching parallel applications under tools and debuggers:
4 indirect and direct. In the indirect launching method (Section 18.2.2, the tool is attached to the
5 starter. In the direct launching method (Section 18.2.1, the tool takes the place of the starter. PMIx
6 also supports attaching to already running programs via the *Process Acquisition* interfaces
7 (Section B.2.1.4).

8 B.2.1.3 Process Synchronization

9 Process Synchronization is a technique tools use to start the processes of a parallel application such
10 that the tools can still attach to the process early in its lifetime. Said another way, the tool must be
11 able to start the application processes without them “running away” from the tool. In the case of
12 MPI (Version 3.1 [4] or the MPI World Process in future versions), this means stopping the
13 applications processes before they return from `MPI_Init` or `MPI_Init_thread`.

14 B.2.1.4 Process Acquisition

15 Process Acquisition is a technique tools use to locate all of the processes, local and remote, of a
16 given parallel application. This typically boils down to collecting the following information for
17 every process in the parallel application: the hostname or IP of the machine running the process,
18 the executable name, and the process ID.

19 B.2.2 Use Case Details

20 B.2.2.1 Direct-Launch Debugger Tool

21 PMIx can support the tool itself using the PMIx spawn options to control the app’s startup,
22 including directing the RM/application as to when to block and wait for tool attachment, or
23 stipulating that an interceptor library be preloaded. However, this means that the user is restricted to
24 whatever command line options the tool vendor has provided for operations such as process
25 placement and binding, which places a significant burden on the tool vendor. An example might
26 look like the following: `dbgr -n 3 ./myapp`.

27 Assuming it is supported, co-launch of debugger daemons in this use-case is supported by adding a
28 `pmix_app_t` to the `PMIx_Spawn` command, indicating that the resulting processes are
29 debugger daemons by setting the `PMIX_DEBUGGER_DAEMONS` attribute.

30 Related Interfaces

31 *PMIx v2.0*

[PMIx_tool_init](#)

C

```
32 pmix_status_t  
33 PMIx_tool_init(pmix_proc_t *proc,  
34               pmix_info_t info[], size_t ninfo);
```




Figure B.1.: Interaction diagram showing an example of the Direct Launch mechanism

```

1 PMIx v2.0 PMIx_Register_event_handler
2
3 pmix_status_t
4 PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,
5 pmix_info_t info[], size_t ninfo,
6 pmix_notification_fn_t evhdlr,
7 pmix_hdlr_reg_cbfunc_t cbfunc,
8 void *cbdata);
9
10 PMIx v4.0 PMIx_Query_info
11
12 pmix_status_t
13 PMIx_Query_info(pmix_query_t queries[], size_t nqueries,
14 pmix_info_t *info[], size_t *ninfo);
15
16 PMIx v1.0 PMIx_Spawn
17
18 pmix_status_t
19 PMIx_Spawn(const pmix_info_t job_info[], size_t ninfo,
20 const pmix_app_t apps[], size_t napps,
21 char nspace[])
22
23 PMIx v1.0 PMIx_Get
24
25 pmix_status_t
26 PMIx_Get(const pmix_proc_t *proc, const pmix_key_t key,
27 const pmix_info_t info[], size_t ninfo,
28 pmix_value_t **val);
29
30 PMIx v2.0 PMIx_Notify_event

```



```

1 pmix_status_t
2 PMIx_Notify_event(pmix_status_t status,
3                 const pmix_proc_t *source,
4                 pmix_data_range_t range,
5                 pmix_info_t info[], size_t ninfo,
6                 pmix_op_cbfunc_t cbfunc, void *cbdata);

```

Related Attributes

PMIX_QUERY_SPAWN_SUPPORT "pmix.qry.spawn" (bool)

Return a comma-delimited list of supported spawn attributes. NO QUALIFIERS.

PMIX_QUERY_DEBUG_SUPPORT "pmix.qry.debug" (bool)

Return a comma-delimited list of supported debug attributes. NO QUALIFIERS.

PMIX_DEBUG_STOP_IN_INIT "pmix.dbg.init" (bool)

Included in either the `pmix_info_t` array in a `pmix_app_t` description (if the directive applies only to that application) or in the `job_info` array if it applies to all applications in the given spawn request. Indicates that the specified application is being spawned under a debugger. The PMIx client library in each resulting application process shall notify its PMIx server that it is pausing and then pause during `PMIx_Init` of the spawned processes until either released by debugger modification of an appropriate variable or receipt of the `PMIX_DEBUGGER_RELEASE` event. The launcher (RM or IL) is responsible for generating the `PMIX_READY_FOR_DEBUG` event (stipulating a breakpoint of `pmix-init`) when all processes have reached the pause point.

PMIX_DEBUG_STOP_ON_EXEC "pmix.dbg.exec" (bool)

Included in either the `pmix_info_t` array in a `pmix_app_t` description (if the directive applies only to that application) or in the `job_info` array if it applies to all applications in the given spawn request. Indicates that the application is being spawned under a debugger, and that the local launch agent is to pause the resulting application processes on first instruction for debugger attach. The launcher (RM or IL) is to generate the `PMIX_LAUNCH_COMPLETE` event when all processes are stopped at the exec point.

PMIX_DEBUG_DAEMONS_PER_PROC "pmix.dbg.dpproc" (uint16_t)

Number of debugger daemons to be spawned per application process. The launcher is to pass the identifier of the namespace to be debugged by including the `PMIX_DEBUG_TARGET` attribute in the daemon's job-level information. The debugger daemons spawned on a given node are responsible for self-determining their specific target process(es) - e.g., by referencing their own `PMIX_LOCAL_RANK` in the daemon debugger job versus the corresponding `PMIX_LOCAL_RANK` of the target processes on the node.

PMIX_DEBUG_DAEMONS_PER_NODE "pmix.dbg.dpnd" (uint16_t)

1 Number of debugger daemons to be spawned on each node where the target job is executing.
2 The launcher is to pass the identifier of the namespace to be debugged by including the
3 **PMIX_DEBUG_TARGET** attribute in the daemon's job-level information. The debugger
4 daemons spawned on a given node are responsible for self-determining their specific target
5 process(es) - e.g., by referencing their own **PMIX_LOCAL_RANK** in the daemon debugger
6 job versus the corresponding **PMIX_LOCAL_RANK** of the target processes on the node.

7 **PMIX_COSPAWN_APP** "pmix.cospawn" (bool)

8 Designated application is to be spawned as a disconnected job - i.e., the launcher shall not
9 include the application in any of the job-level values (e.g., **PMIX_RANK** within the job)
10 provided to any other application process generated by the same spawn request. Typically
11 used to cospawn debugger daemons alongside an application.

12 **PMIX_MAPBY** "pmix.mapby" (char*)

13 Process mapping policy - when accessed using **PMIx_Get**, use the
14 **PMIX_RANK_WILDCARD** value for the rank to discover the mapping policy used for the
15 provided namespace. Supported values are launcher specific.

16 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)

17 Requests that the ability to forward the **stdout** of the spawned processes be maintained.
18 The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and
19 other options for delivery of the forwarded output.

20 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)

21 Requests that the ability to forward the **stderr** of the spawned processes be maintained.
22 The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and
23 other options for delivery of the forwarded output.

24 **PMIX_NOTIFY_COMPLETION** "pmix.notecomp" (bool)

25 Requests that the launcher generate the **PMIX_EVENT_JOB_END** event for normal or
26 abnormal termination of the spawned job. The event shall include the returned status code
27 (**PMIX_JOB_TERM_STATUS**) for the corresponding job; the identity (**PMIX_PROCID**)
28 and exit status (**PMIX_EXIT_CODE**) of the first failed process, if applicable; and a
29 **PMIX_EVENT_TIMESTAMP** indicating the time the termination occurred. Note that the
30 requester must register for the event or capture and process it within a default event handler.

31 **PMIX_SETUP_APP_ENVARS** "pmix.setup.env" (bool)

32 Harvest and include relevant environmental variables.

33 **PMIX_EVENT_AFFECTED_PROC** "pmix.evproc" (pmix_proc_t)

34 The single process that was affected.

35 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)

36 Included in the **pmix_info_t** array of a **pmix_app_t**, this attribute declares that the
37 application consists of debugger daemons and shall be governed accordingly. If used as the
38 sole **pmix_app_t** in a **PMIx_Spawn** request, then the **PMIX_DEBUG_TARGET** attribute
39 must also be provided (in either the *job_info* or in the *info* array of the **pmix_app_t**) to

1 identify the namespace to be debugged so that the launcher can determine where to place the
2 spawned daemons. If neither `PMIX_DEBUG_DAEMONS_PER_PROC` nor
3 `PMIX_DEBUG_DAEMONS_PER_NODE` is specified, then the launcher shall default to a
4 placement policy of one daemon per process in the target job.

5 `PMIX_DEBUG_TARGET` "pmix.dbg.tgt" (`pmix_proc_t*`)

6 Identifier of process(es) to be debugged - a rank of `PMIX_RANK_WILDCARD` indicates that
7 all processes in the specified namespace are to be included.

8 `PMIX_DEBUG_STOP_IN_APP` "pmix.dbg.notify" (`varies`)

9 Direct specified ranks to stop at application-specific point and notify they are
10 ready-to-debug. The attribute's value can be any of three data types:

- 11 • `bool` - true indicating all ranks
- 12 • `pmix_rank_t` - the rank of one proc, or `PMIX_RANK_WILDCARD` for all
- 13 • a `pmix_data_array_t` if an array of individual processes are specified

14 The resulting application processes are to notify their server (by generating the
15 `PMIX_READY_FOR_DEBUG` event) when they reach some application-determined location
16 - the event shall include the `PMIX_BREAKPOINT` attribute indicating where the application
17 has stopped. The application shall pause at that point until released by debugger
18 modification of an appropriate variable. The launcher (RM or IL) is responsible for
19 generating the `PMIX_READY_FOR_DEBUG` event when all processes have indicated they
20 are at the pause point.

21 `PMIX_QUERY_LOCAL_PROC_TABLE` "pmix.qry.lptable" (`char*`)

22 Returns a (`pmix_data_array_t`) array of `pmix_proc_info_t`, one entry for each
23 process in the specified namespace executing on the same node as the requester, ordered by
24 process job rank. REQUIRED QUALIFIER: `PMIX_NAMESPACE` indicating the namespace
25 whose local process table is being queried. OPTIONAL QUALIFIER: `PMIX_HOSTNAME`
26 indicating the host whose local process table is being queried. By default, the query assumes
27 that the host upon which the request was made is to be used.

28 Related Constants

29 `PMIX_DEBUG_WAITING_FOR_NOTIFY`

30 `PMIX_DEBUGGER_RELEASE`

31 B.2.2.2 Indirect-Launch Debugger Tool

32 Executing a program under a tool using an intermediate launcher such as `mpiexec` can also be
33 made possible. This requires some degree of coordination between the tool and the launcher.
34 Ultimately, it is the launcher that is going to launch the application, and the tool must somehow
35 inform the launcher (and the application) that this is being done in a debug session so that the
36 application knows to "block" until the tool attaches to it.

37 In this operational mode, the user invokes a tool (typically on a non-compute, or "head", node) that
38 in turn uses `mpiexec` to launch their application – a typical command line might look like the
39 following: `dbgr -dbgoption mpiexec -n 32 ./myapp`.

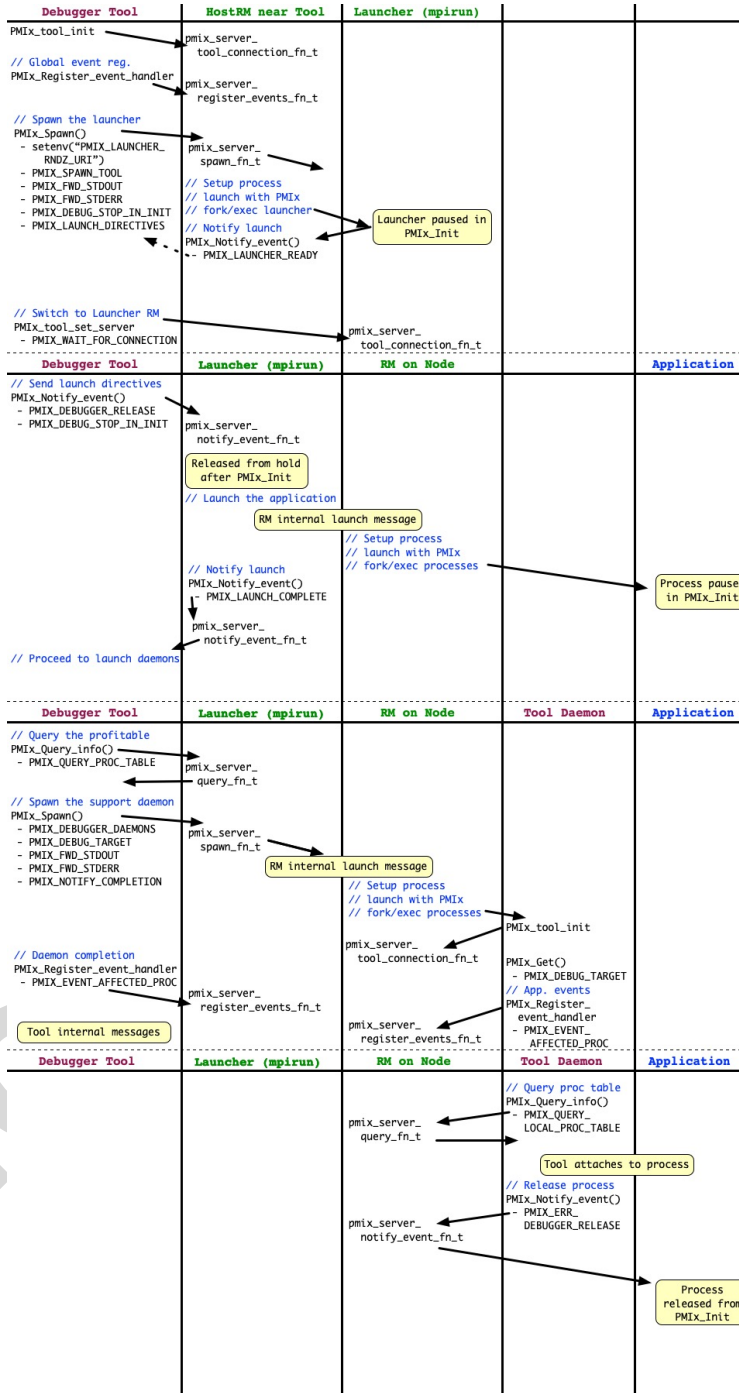


Figure B.2.: Interaction diagram showing an example of the Indirect Launch mechanism

1 **Related Interfaces**

2 *PMIx v2.0* **PMIx_tool_init**



```
3 pmix_status_t  
4 PMIx_tool_init(pmix_proc_t *proc,  
5 pmix_info_t info[], size_t ninfo);
```



6 *PMIx v2.0* **PMIx_Register_event_handler**



```
7 pmix_status_t  
8 PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,  
9 pmix_info_t info[], size_t ninfo,  
10 pmix_notification_fn_t evhdlr,  
11 pmix_hdlr_reg_cbfunc_t cbfunc,  
12 void *cbdata);
```



13 *PMIx v1.0* **PMIx_Spawn**



```
14 pmix_status_t  
15 PMIx_Spawn(const pmix_info_t job_info[], size_t ninfo,  
16 const pmix_app_t apps[], size_t napps,  
17 char nspace[])
```



18 *PMIx v2.0* **PMIx_Notify_event**



```
19 pmix_status_t  
20 PMIx_Notify_event(pmix_status_t status,  
21 const pmix_proc_t *source,  
22 pmix_data_range_t range,  
23 pmix_info_t info[], size_t ninfo,  
24 pmix_op_cbfunc_t cbfunc, void *cbdata);
```



25 *PMIx v4.0* **PMIx_tool_attach_to_server**

```

1  pmix_status_t
2  PMIx_tool_attach_to_server(pmix_proc_t *proc,
3                             pmix_proc_t *server,
4                             pmix_info_t info[],
5                             size_t ninfo);

```

6 *PMIx v4.0* **PMIx_Query_info**

```

7  pmix_status_t
8  PMIx_Query_info(pmix_query_t queries[], size_t nqueries,
9                 pmix_info_t *info[], size_t *ninfo);

```

10 *PMIx v1.0* **PMIx_Get**

```

11 pmix_status_t
12 PMIx_Get(const pmix_proc_t *proc, const pmix_key_t key,
13          const pmix_info_t info[], size_t ninfo,
14          pmix_value_t **val);

```

15 Related Attributes

16 **PMIX_LAUNCH_DIRECTIVES** "pmix.lnch.dirs" (pmix_data_array_t*)

17 Array of **pmix_info_t** containing directives for the launcher - a convenience attribute for
18 retrieving all directives with a single call to **PMIx_Get**.

19 **PMIX_SPAWN_TOOL** "pmix.spwn.tool" (bool)

20 Indicate that the job being spawned is a tool.

21 **PMIX_COSPAWN_APP** "pmix.cospawn" (bool)

22 Designated application is to be spawned as a disconnected job - i.e., the launcher shall not
23 include the application in any of the job-level values (e.g., **PMIX_RANK** within the job)
24 provided to any other application process generated by the same spawn request. Typically
25 used to cospawn debugger daemons alongside an application.

26 **PMIX_FWD_STDOUT** "pmix.fwd.stdout" (bool)

27 Requests that the ability to forward the **stdout** of the spawned processes be maintained.
28 The requester will issue a call to **PMIx_IOF_pull** to specify the callback function and
29 other options for delivery of the forwarded output.

30 **PMIX_FWD_STDERR** "pmix.fwd.stderr" (bool)

1 Requests that the ability to forward the `stderr` of the spawned processes be maintained.
2 The requester will issue a call to `PMIx_IOF_pull` to specify the callback function and
3 other options for delivery of the forwarded output.

4 `PMIX_SETUP_APP_ENVARS` "pmix.setup.env" (bool)

5 Harvest and include relevant environmental variables.

6 `PMIX_DEBUG_STOP_IN_INIT` "pmix.dbg.init" (bool)

7 Included in either the `pmix_info_t` array in a `pmix_app_t` description (if the directive
8 applies only to that application) or in the `job_info` array if it applies to all applications in the
9 given spawn request. Indicates that the specified application is being spawned under a
10 debugger. The PMIx client library in each resulting application process shall notify its PMIx
11 server that it is pausing and then pause during `PMIx_Init` of the spawned processes until
12 either released by debugger modification of an appropriate variable or receipt of the
13 `PMIX_DEBUGGER_RELEASE` event. The launcher (RM or IL) is responsible for generating
14 the `PMIX_READY_FOR_DEBUG` event (stipulating a breakpoint of `pmix-init`) when all
15 processes have reached the pause point.

16 `PMIX_DEBUG_STOP_ON_EXEC` "pmix.dbg.exec" (bool)

17 Included in either the `pmix_info_t` array in a `pmix_app_t` description (if the directive
18 applies only to that application) or in the `job_info` array if it applies to all applications in the
19 given spawn request. Indicates that the application is being spawned under a debugger, and
20 that the local launch agent is to pause the resulting application processes on first instruction
21 for debugger attach. The launcher (RM or IL) is to generate the
22 `PMIX_LAUNCH_COMPLETE` event when all processes are stopped at the exec point.

23 `PMIX_DEBUG_DAEMONS_PER_PROC` "pmix.dbg.dpproc" (uint16_t)

24 Number of debugger daemons to be spawned per application process. The launcher is to pass
25 the identifier of the namespace to be debugged by including the `PMIX_DEBUG_TARGET`
26 attribute in the daemon's job-level information. The debugger daemons spawned on a given
27 node are responsible for self-determining their specific target process(es) - e.g., by
28 referencing their own `PMIX_LOCAL_RANK` in the daemon debugger job versus the
29 corresponding `PMIX_LOCAL_RANK` of the target processes on the node.

30 `PMIX_DEBUG_DAEMONS_PER_NODE` "pmix.dbg.dpnd" (uint16_t)

31 Number of debugger daemons to be spawned on each node where the target job is executing.
32 The launcher is to pass the identifier of the namespace to be debugged by including the
33 `PMIX_DEBUG_TARGET` attribute in the daemon's job-level information. The debugger
34 daemons spawned on a given node are responsible for self-determining their specific target
35 process(es) - e.g., by referencing their own `PMIX_LOCAL_RANK` in the daemon debugger
36 job versus the corresponding `PMIX_LOCAL_RANK` of the target processes on the node.

37 `PMIX_MAPBY` "pmix.mapby" (char*)

38 Process mapping policy - when accessed using `PMIx_Get`, use the
39 `PMIX_RANK_WILDCARD` value for the rank to discover the mapping policy used for the
40 provided namespace. Supported values are launcher specific.

1 **PMIX_QUERY_PROC_TABLE** "pmix.qry.phtable" (char*)
 2 Returns a (**pmix_data_array_t**) array of **pmix_proc_info_t**, one entry for each
 3 process in the specified namespace, ordered by process job rank. REQUIRED QUALIFIER:
 4 **PMIX_NAMESPACE** indicating the namespace whose process table is being queried.

5 **PMIX_QUERY_LOCAL_PROC_TABLE** "pmix.qry.lhtable" (char*)
 6 Returns a (**pmix_data_array_t**) array of **pmix_proc_info_t**, one entry for each
 7 process in the specified namespace executing on the same node as the requester, ordered by
 8 process job rank. REQUIRED QUALIFIER: **PMIX_NAMESPACE** indicating the namespace
 9 whose local process table is being queried. OPTIONAL QUALIFIER: **PMIX_HOSTNAME**
 10 indicating the host whose local process table is being queried. By default, the query assumes
 11 that the host upon which the request was made is to be used.

12 **PMIX_DEBUGGER_DAEMONS** "pmix.debugger" (bool)
 13 Included in the **pmix_info_t** array of a **pmix_app_t**, this attribute declares that the
 14 application consists of debugger daemons and shall be governed accordingly. If used as the
 15 sole **pmix_app_t** in a **PMIX_Spawn** request, then the **PMIX_DEBUG_TARGET** attribute
 16 must also be provided (in either the *job_info* or in the *info* array of the **pmix_app_t**) to
 17 identify the namespace to be debugged so that the launcher can determine where to place the
 18 spawned daemons. If neither **PMIX_DEBUG_DAEMONS_PER_PROC** nor
 19 **PMIX_DEBUG_DAEMONS_PER_NODE** is specified, then the launcher shall default to a
 20 placement policy of one daemon per process in the target job.

21 **PMIX_NOTIFY_COMPLETION** "pmix.notecomp" (bool)
 22 Requests that the launcher generate the **PMIX_EVENT_JOB_END** event for normal or
 23 abnormal termination of the spawned job. The event shall include the returned status code
 24 (**PMIX_JOB_TERM_STATUS**) for the corresponding job; the identity (**PMIX_PROCID**)
 25 and exit status (**PMIX_EXIT_CODE**) of the first failed process, if applicable; and a
 26 **PMIX_EVENT_TIMESTAMP** indicating the time the termination occurred. Note that the
 27 requester must register for the event or capture and process it within a default event handler.

28 **PMIX_DEBUG_TARGET** "pmix.dbg.tgt" (**pmix_proc_t***)
 29 Identifier of process(es) to be debugged - a rank of **PMIX_RANK_WILDCARD** indicates that
 30 all processes in the specified namespace are to be included.

31 **PMIX_WAIT_FOR_CONNECTION** "pmix.wait.conn" (bool)
 32 Wait until the specified process has connected to the requesting tool or server, or the
 33 operation times out (if the **PMIX_TIMEOUT** directive is included in the request).

34 Related Constants

35 **PMIX_LAUNCHER_READY**
 36 **PMIX_LAUNCH_COMPLETE**
 37 **PMIX_DEBUG_WAITING_FOR_NOTIFY**
 38 **PMIX_DEBUGGER_RELEASE**
 39 **PMIX_LAUNCHER_RNDZ_URI**

1 B.2.2.3 Attaching to a Running Job

2 PMIx supports attaching to an already running parallel job in two ways. In the first way, the main
3 process of a tool calls `PMIx_Query_info` with the `PMIX_QUERY_PROC_TABLE` attribute.
4 This returns an array of structs containing the information required for [process acquisition](#). This
5 includes remote hostnames, executable names, and process IDs. In the second way, every tool
6 daemon calls `PMIx_Query_info` with the `PMIX_QUERY_LOCAL_PROC_TABLE` attribute.
7 This returns a similar array of structs but only for processes on the same node.

8 An example of this use-case may look like the following: `mpirun -n 32 ./myApp &&`
9 `dbgr attach $!`.

10 *PMIx v2.0*

`PMIx_tool_init`

```
▼ C ▼  
11 pmix_status_t  
12 PMIx_tool_init(pmix_proc_t *proc,  
13                pmix_info_t info[], size_t ninfo);  
▲ C ▲
```

14 *PMIx v2.0*

`PMIx_Register_event_handler`

```
▼ C ▼  
15 pmix_status_t  
16 PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,  
17                             pmix_info_t info[], size_t ninfo,  
18                             pmix_notification_fn_t evhdlr,  
19                             pmix_hdlr_reg_cbfunc_t cbfunc,  
20                             void *cbdata);  
▲ C ▲
```

21 *PMIx v4.0*

`PMIx_Query_info`

```
▼ C ▼  
22 pmix_status_t  
23 PMIx_Query_info(pmix_query_t queries[], size_t nqueries,  
24                 pmix_info_t *info[], size_t *ninfo);  
▲ C ▲
```

25 *PMIx v1.0*

`PMIx_Spawn`




Figure B.3.: Interaction diagram showing an example of the attaching to a running job

C

```
1 pmix_status_t
2 PMix_Spawn(const pmix_info_t job_info[], size_t ninfo,
3             const pmix_app_t apps[], size_t napps,
4             char nspace[])
```

C

Related Attributes

PMIX_QUERY_PROC_TABLE "pmix.qry.ptable" (char*)

Returns a ([pmix_data_array_t](#)) array of [pmix_proc_info_t](#), one entry for each process in the specified namespace, ordered by process job rank. REQUIRED QUALIFIER: [PMIX_NAMESPACE](#) indicating the namespace whose process table is being queried.

PMIX_DEBUGGER_DAEMONS "pmix.debugger" (bool)

Included in the [pmix_info_t](#) array of a [pmix_app_t](#), this attribute declares that the application consists of debugger daemons and shall be governed accordingly. If used as the sole [pmix_app_t](#) in a [PMix_Spawn](#) request, then the [PMIX_DEBUG_TARGET](#) attribute must also be provided (in either the *job_info* or in the *info* array of the [pmix_app_t](#)) to identify the namespace to be debugged so that the launcher can determine where to place the spawned daemons. If neither [PMIX_DEBUG_DAEMONS_PER_PROC](#) nor [PMIX_DEBUG_DAEMONS_PER_NODE](#) is specified, then the launcher shall default to a placement policy of one daemon per process in the target job.

PMIX_DEBUG_TARGET "pmix.dbg.tgt" ([pmix_proc_t*](#))

Identifier of process(es) to be debugged - a rank of [PMIX_RANK_WILDCARD](#) indicates that all processes in the specified namespace are to be included.

PMIX_DEBUG_DAEMONS_PER_PROC "pmix.dbg.dpproc" ([uint16_t](#))

Number of debugger daemons to be spawned per application process. The launcher is to pass the identifier of the namespace to be debugged by including the [PMIX_DEBUG_TARGET](#) attribute in the daemon's job-level information. The debugger daemons spawned on a given node are responsible for self-determining their specific target process(es) - e.g., by referencing their own [PMIX_LOCAL_RANK](#) in the daemon debugger job versus the corresponding [PMIX_LOCAL_RANK](#) of the target processes on the node.

PMIX_DEBUG_DAEMONS_PER_NODE "pmix.dbg.dpnd" ([uint16_t](#))

Number of debugger daemons to be spawned on each node where the target job is executing. The launcher is to pass the identifier of the namespace to be debugged by including the [PMIX_DEBUG_TARGET](#) attribute in the daemon's job-level information. The debugger daemons spawned on a given node are responsible for self-determining their specific target process(es) - e.g., by referencing their own [PMIX_LOCAL_RANK](#) in the daemon debugger job versus the corresponding [PMIX_LOCAL_RANK](#) of the target processes on the node.

PMIX_MAPBY "pmix.mapby" (char*)

1 Process mapping policy - when accessed using `PMIx_Get`, use the
2 `PMIX_RANK_WILDCARD` value for the rank to discover the mapping policy used for the
3 provided namespace. Supported values are launcher specific.

4 `PMIX_FWD_STDOUT` "`pmix.fwd.stdout`" (bool)

5 Requests that the ability to forward the `stdout` of the spawned processes be maintained.
6 The requester will issue a call to `PMIx_IOF_pull` to specify the callback function and
7 other options for delivery of the forwarded output.

8 `PMIX_FWD_STDERR` "`pmix.fwd.stderr`" (bool)

9 Requests that the ability to forward the `stderr` of the spawned processes be maintained.
10 The requester will issue a call to `PMIx_IOF_pull` to specify the callback function and
11 other options for delivery of the forwarded output.

12 `PMIX_NOTIFY_COMPLETION` "`pmix.notecomp`" (bool)

13 Requests that the launcher generate the `PMIX_EVENT_JOB_END` event for normal or
14 abnormal termination of the spawned job. The event shall include the returned status code
15 (`PMIX_JOB_TERM_STATUS`) for the corresponding job; the identity (`PMIX_PROCID`)
16 and exit status (`PMIX_EXIT_CODE`) of the first failed process, if applicable; and a
17 `PMIX_EVENT_TIMESTAMP` indicating the time the termination occurred. Note that the
18 requester must register for the event or capture and process it within a default event handler.

19 `PMIX_REQUESTOR_IS_TOOL` "`pmix.req.tool`" (bool)

20 The requesting process is a PMIx tool.

21 `PMIX_QUERY_NAMESPACES` "`pmix.qry.ns`" (char*)

22 Request a comma-delimited list of active namespaces. NO QUALIFIERS.

23 B.2.2.4 Tool Interaction with RM

24 Tools can benefit from a mechanism by which they may interact with a local PMIx server that has
25 opted to accept such connections along with support for tool connections to system-level PMIx
26 servers, and a logging feature. To add support for tool connections to a specified system-level,
27 PMIx server environments could choose to launch a set of PMIx servers to support a given
28 allocation - these servers will (if so instructed) provide a tool rendezvous point that is tagged with
29 their pid and typically placed in an allocation-specific temporary directory to allow for possible
30 multi-tenancy scenarios. Supporting such operations requires that a system-level PMIx connection
31 be provided which is not associated with a specific user or allocation. A new key has been added to
32 direct the PMIx server to expose a rendezvous point specifically for this purpose.

33 *PMIx v2.0* `PMIx_Query_info_nb`

```

1      pmix_status_t
2      PMIx_Query_info_nb(pmix_query_t queries[], size_t nqueries,
3                          pmix_info_cbfunc_t cbfunc, void *cbdata);
4 PMIx v2.0  PMIx_Register_event_handler
5
6      pmix_status_t
7      PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,
8                                  pmix_info_t info[], size_t ninfo,
9                                  pmix_notification_fn_t evhdlr,
10                                 pmix_hdlr_reg_cbfunc_t cbfunc,
11                                 void *cbdata);
12 PMIx v2.0  PMIx_Deregister_event_handler
13
14      pmix_status_t
15      PMIx_Deregister_event_handler(size_t evhdlr_ref,
16                                     pmix_op_cbfunc_t cbfunc,
17                                     void *cbdata);
18 PMIx v2.0  PMIx_Notify_event
19
20      pmix_status_t
21      PMIx_Notify_event(pmix_status_t status,
22                        const pmix_proc_t *source,
23                        pmix_data_range_t range,
24                        pmix_info_t info[], size_t ninfo,
25                        pmix_op_cbfunc_t cbfunc, void *cbdata);
26 PMIx v1.0  PMIx_server_init

```

```

1  pmix_status_t
2  PMIx_server_init(pmix_server_module_t *module,
3                  pmix_info_t info[], size_t ninfo);

```

B.2.2.5 Environmental Parameter Directives for Applications and Launchers

It is sometimes desirable or required that standard environmental variables (e.g., `PATH`, `LD_LIBRARY_PATH`, `LD_PRELOAD`) be modified prior to executing an application binary or a starter such as `mpixexec` - this is particularly true when tools/debuggers are used to start the application.

Related Interfaces

PMIx v1.0 [PMIx_Spawn](#)

```

12 pmix_status_t
13 PMIx_Spawn(const pmix_info_t job_info[], size_t ninfo,
14            const pmix_app_t apps[], size_t napps,
15            char nspace[])

```

Related Structs

[pmix_env_var_t](#)

Related Attributes

[PMIX_SET_ENVAR](#) "pmix.envar.set" (pmix_env_var_t*)

Set the envar to the given value, overwriting any pre-existing one

[PMIX_ADD_ENVAR](#) "pmix.envar.add" (pmix_env_var_t*)

Add the environment variable, but do not overwrite any pre-existing one

[PMIX_UNSET_ENVAR](#) "pmix.envar.unset" (char*)

Unset the environment variable specified in the string.

[PMIX_PREPEND_ENVAR](#) "pmix.envar.prepnd" (pmix_env_var_t*)

Prepend the given value to the specified environmental value using the given separator character, creating the variable if it doesn't already exist

[PMIX_APPEND_ENVAR](#) "pmix.envar.appnd" (pmix_env_var_t*)

Append the given value to the specified environmental value using the given separator character, creating the variable if it doesn't already exist

Resource managers and launchers must scan for relevant directives, modifying environmental parameters as directed. Directives are to be processed in the order in which they were given, starting with job-level directives (applied to each app) followed by app-level directives.

1 **B.3 Hybrid Applications**

2 **B.3.1 Use Case Summary**

3 Hybrid applications (i.e., applications that utilize more than one programming model or runtime
4 system, such as an application using MPI that also uses OpenMP or UPS) are growing in
5 popularity, especially as processors with increasingly large numbers of cores and/or hardware
6 threads proliferate. Unfortunately, the various corresponding runtime systems currently operate
7 under the assumption that they alone control execution. This leads to conflicts in hybrid
8 applications. Deadlock of parallel applications can occur when one runtime system prevents the
9 other from making progress due to lack of coordination between them [3]. Sub-optimal
10 performance can also occur due to uncoordinated division of hardware resources between the
11 runtime systems implementing the different programming models or systems [5, 6]. This use-case
12 offers potential solutions to this problem by providing a pathway for parallel runtime systems to
13 coordinate their actions.

14 **B.3.2 Use Case Details**

15 **B.3.2.1 Identifying Active Parallel Runtime Systems**

16 The current state-of-the-practice for concurrently used runtime systems in a single application to
17 detect one another is via set environment variables. For example, some OpenMP implementations
18 look for environment variables to indicate that an MPI library is active. Unfortunately, this
19 technique is not completely reliable as environment variables change over time and with new
20 software versions, and this detection is implementation specific. Also, the fact that an environment
21 variable is present doesn't guarantee that a particular runtime system is in active use since Resource
22 Managers routinely set environment variables "just in case" the application needs them. PMIx
23 provides a reliable mechanism by which each library can determine that another runtime library is
24 in operation.

25 When initializing PMIx, runtime libraries implementing a parallel programming model can register
26 themselves, including their name, the library version, the version of the API they implement, and
27 the threading model. This information is then cached locally and can then be read asynchronously
28 by other runtime systems using PMIx's Event Notification system.

29 This initialization mechanism also allows runtime libraries to share knowledge of each other's
30 resources and intended resource utilization. For example, if an OpenMP implementation knows
31 which hardware threads an MPI library is using it could potentially avoid core and cache contention.

32 **Code Example**

```
1 pmix_proc_t myproc;  
2 pmix_info_t *info;  
3 volatile bool wearedone = false;  
4
```


```

5 PMIX_INFO_CREATE(info, 4);
6 PMIX_INFO_LOAD(&info[0], PMIX_PROGRAMMING_MODEL, "MPI", PMIX_STRING);
7 PMIX_INFO_LOAD(&info[1], PMIX_MODEL_LIBRARY_NAME, "FooMPI",
  ↪ PMIX_STRING);
8 PMIX_INFO_LOAD(&info[2], PMIX_MODEL_LIBRARY_VERSION, "1.0.0",
  ↪ PMIX_STRING);
9 PMIX_INFO_LOAD(&info[3], PMIX_THREADING_MODEL, "posix", PMIX_STRING);
10 pmix_status_t rc = PMIx_Init(&myproc, info, 4);
11 PMIX_INFO_FREE(info, 4);

```

1 Related Interfaces

2 *PMIx v1.2* **PMIx_Init**



6 Related Attributes

7 **PMIX_PROGRAMMING_MODEL** "pmix.pgm.model" (char*)

8 Programming model being initialized (e.g., "MPI" or "OpenMP").

9 **PMIX_MODEL_LIBRARY_NAME** "pmix.mdl.name" (char*)

10 Programming model implementation ID (e.g., "OpenMPI" or "MPICH").

11 **PMIX_MODEL_LIBRARY_VERSION** "pmix.mdl.vrs" (char*)

12 Programming model version string (e.g., "2.1.1").

13 **PMIX_THREADING_MODEL** "pmix.threads" (char*)

14 Threading model used (e.g., "pthreads").

15 **PMIX_MODEL_NUM_THREADS** "pmix.mdl.nthrds" (uint64_t)

16 Number of active threads being used by the model.

17 **PMIX_MODEL_NUM_CPUS** "pmix.mdl.ncpu" (uint64_t)

18 Number of cpus being used by the model.

19 **PMIX_MODEL_CPU_TYPE** "pmix.mdl.cputype" (char*)

20 Granularity - "hwthread", "core", etc.

21 **PMIX_MODEL_PHASE_NAME** "pmix.mdl.phase" (char*)

22 User-assigned name for a phase in the application execution (e.g., "cfd reduction").

23 **PMIX_MODEL_PHASE_TYPE** "pmix.mdl.ptype" (char*)

24 Type of phase being executed (e.g., "matrix multiply").

1 **PMIX_MODEL_AFFINITY_POLICY** "pmix.mdl.tap" (char*)
2 Thread affinity policy - e.g.: "master" (thread co-located with master thread), "close" (thread
3 located on cpu close to master thread), "spread" (threads load-balanced across available
4 cpus).

5 B.3.2.2 Coordinating at Runtime

6 The PMIx Event Notification system provides a mechanism by which the resource manager can
7 communicate system events to applications, thus providing applications with an opportunity to
8 generate an appropriate response. Hybrid applications can leverage these events for cross-library
9 coordination.

10 Runtime libraries can access the information provided by other runtime libraries during their
11 initialization using the event notification system. In this case, runtime libraries should register a
12 callback for the **PMIX_MODEL_DECLARED** event.

13 Applications, runtime libraries, and resource managers can also use the PMIx event notification
14 system to communicate dynamic information, such as entering a new application phase
15 (**PMIX_MODEL_PHASE_NAME**) or a change in resources used (**PMIX_MODEL_RESOURCES**).
16 This dynamic information can be broadcast using the **PMIx_Notify_event** function. Runtime
17 libraries can register callback functions to run when these events occur using
18 **PMIx_Register_event_handler**.

19 Code Example

20 Registering a callback to run when another runtime library initializes:

```
1 static void model_declared_cb(size_t evhdlr_registration_id,  
2                               pmix_status_t status, const pmix_proc_t  
3   ↪ *source,  
4                               pmix_info_t info[], size_t ninfo,  
5                               pmix_info_t results[], size_t nresults,  
6   ↪ cbfunc,  
7                               void *cbdata) {  
8     printf("Entered %s\n", __FUNCTION__);  
9     int n;  
10    for (n = 0; n < ninfo; n++) {  
11      if (PMIX_CHECK_KEY(&info[n], PMIX_PROGRAMMING_MODEL) &&  
12          strcmp(info[n].value.data.string, "MPI") == 0) {  
13        /* ignore our own declaration */  
14        break;  
15      } else {  
16        /* actions to perform when another model registers */  
17      }  
18    }  
19    if (NULL != cbfunc) {  
20      /* tell the event handler that we are only a partial step */
```

```

20     cbfunc(PMIX_EVENT_PARTIAL_ACTION_TAKEN, NULL, 0, NULL, NULL,
↳     cbdata);
21     }
22 }
23
24     pmix_status_t code = PMIX_MODEL_DECLARED;
25     rc = PMIx_Register_event_handler(&code, 1, NULL, 0, model_declared_cb,
↳     NULL, NULL);

```

1 Notifying an event:

```

1 PMIX_INFO_CREATE(info, 1);
2 PMIX_INFO_LOAD(&info[0], PMIX_EVENT_NON_DEFAULT, NULL, PMIX_BOOL);
3 rc = PMIx_Notify_event(PMIX_OPENMP_PARALLEL_ENTERED, &myproc,
↳ PMIX_RANGE_PROC_LOCAL, info, 1, notify_complete, (void*)&wearedone);

```

2 Related Interfaces

3 *PMIx v2.0* [PMIx_Notify_event](#)

```

↳-----C-----
4     pmix_status_t
5     PMIx_Notify_event(pmix_status_t status,
6                     const pmix_proc_t *source,
7                     pmix_data_range_t range,
8                     pmix_info_t info[], size_t ninfo,
9                     pmix_op_cbfunc_t cbfunc, void *cbdata);
↳-----C-----

```

10 *PMIx v2.0* [PMIx_Register_event_handler](#)

```

↳-----C-----
11     pmix_status_t
12     PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,
13                                pmix_info_t info[], size_t ninfo,
14                                pmix_notification_fn_t evhdlr,
15                                pmix_hdlr_reg_cbfunc_t cbfunc,
16                                void *cbdata);
↳-----C-----

```

17 *PMIx v2.0* [pmix_event_notification_cbfunc_fn_t](#)

C

```
1 typedef void (*pmix_event_notification_cbfunc_fn_t)
2     (pmix_status_t status,
3      pmix_info_t *results, size_t nresults,
4      pmix_op_cbfunc_t cbfunc, void *thiscbdata,
5      void *notification_cbdata);
```

C

6 Related Constants

7 [PMIX_MODEL_DECLARED](#)

8 [PMIX_MODEL_RESOURCES](#)

9 [PMIX_OPENMP_PARALLEL_ENTERED](#)

10 [PMIX_OPENMP_PARALLEL_EXITED](#)

11 [PMIX_EVENT_ACTION_COMPLETE](#)

12 B.3.2.3 Coordinating at Runtime with Multiple Event Handlers

13 Coordinating with a threading library such as an OpenMP runtime library creates the need for
14 separate event handlers for threads of the same process. For example in an MPI+OpenMP hybrid
15 application, the MPI main thread and the OpenMP primary thread may both want to be notified
16 anytime an OpenMP thread starts executing in a parallel region. This requires support for multiple
17 threads to potentially register different event handlers against the same status code.

18 Multiple event handlers registered against the same event are processed in a chain-like manner
19 based on the order in which they were registered, as modified by any directives. Registrations
20 against specific event codes are processed first, followed by registrations against multiple event
21 codes and then any default registrations. At each point in the chain, an event handler is called by the
22 PMIx progress thread and given a function to call when that handler has completed its operation.
23 The handler callback notifies PMIx that the handler is done, returning a status code to indicate the
24 result of its work. The results are appended to the array of prior results, with the returned values
25 combined into an array within a single [pmix_info_t](#) as follows:

- 26 • `array[0]`: the event handler name provided at registration (may be an empty field if a string
27 name was not given) will be in the key, with the [pmix_status_t](#) value returned by the handler
- 28 • `array[*]`: the array of results returned by the handler, if any.

29 The current PMIx standard does not actually specify a default ordering for event handlers as they
30 are being registered. However, it does include an inherent ordering for invocation. Specifically,
31 PMIx stipulates that handlers be called in the following categorical order:

- 32 • single status event handlers - handlers that were registered against a single specific status.
- 33 • multi status event handlers - those registered against more than one specific status.
- 34 • default event handlers - those registered against no specific status.

1
2

Code Example

From the OpenMP primary thread:

```
1 static void parallel_region_OMP_cb(size_t evhdlr_registration_id,
2     pmix_status_t status,
3     const pmix_proc_t *source,
4     pmix_info_t info[], size_t ninfo,
5     pmix_info_t results[], size_t
6     ↪ nresults,
7     pmix_event_notification_cbfunc_fn_t
8     ↪ cbfunc,
9     void *cbdata) {
10    printf("Entered %s\n", __FUNCTION__);
11    /* do what we need OpenMP to do on entering a parallel region */
12    if (NULL != cbfunc) {
13        /* tell the event handler that we are only a partial step */
14        cbfunc(PMIX_EVENT_PARTIAL_ACTION_TAKEN, NULL, 0, NULL, NULL,
15        ↪ cbdata);
16    }
17 }
18
19 bool is_true = true;
20 pmix_status_t code = PMIX_OPENMP_PARALLEL_ENTERED;
21 PMIX_INFO_CREATE(info, 2);
22 PMIX_INFO_LOAD(&info[0], PMIX_EVENT_HDLR_NAME, "OpenMP-Primary",
23 ↪ PMIX_STRING);
24 PMIX_INFO_LOAD(&info[1], PMIX_EVENT_HDLR_FIRST, &is_true, PMIX_BOOL);
25 rc = PMIX_Register_event_handler(&code, 1, info, 2,
26 ↪ parallel_region_OMP_cb, NULL, NULL);
27 if (rc < 0)
28     fprintf(stderr, "%s: Failed to register event handler for OpenMP
29 ↪ region entrance\n", __FUNCTION__);
30 PMIX_INFO_FREE(info, 2);
```

3

From the MPI process:

```
1 static void parallel_region_MPI_cb(size_t evhdlr_registration_id,
2     pmix_status_t status,
3     const pmix_proc_t *source,
4     pmix_info_t info[], size_t ninfo,
5     pmix_info_t results[], size_t
6     ↪ nresults,
7     pmix_event_notification_cbfunc_fn_t
8     ↪ cbfunc,
9     void *cbdata) {
10    printf("Entered %s\n", __FUNCTION__);
```

```

9  /* do what we need the MPI library to do on entering a parallel region
↪  */
10  if (NULL != cbfunc) {
11      /* tell the event handler that we are the last step */
12      cbfunc(PMIX_EVENT_ACTION_COMPLETE, NULL, 0, NULL, NULL, cbdata);
13  }
14  }
15
16  pmix_status_t code = PMIX_OPENMP_PARALLEL_ENTERED;
17  PMIX_INFO_CREATE(info, 2);
18  PMIX_INFO_LOAD(&info[0], PMIX_EVENT_HDLR_NAME, "MPI-Thread",
↪  PMIX_STRING);
19  PMIX_INFO_LOAD(&info[1], PMIX_EVENT_HDLR_AFTER, "OpenMP-Primary",
↪  PMIX_STRING);
20  rc = PMIX_Register_event_handler(&code, 1, info, 2,
↪  parallel_region_MPI_cb, NULL, NULL);
21  if (rc < 0)
22      fprintf(stderr, "%s: Failed to register event handler for OpenMP
↪  region entrance\n", __FUNCTION__);
23  PMIX_INFO_FREE(info, 2);

```

1 Related Interfaces

2 *PMIx v2.0* [PMIx_Register_event_handler](#)

↕ C ↕

```

3  pmix_status_t
4  PMIx_Register_event_handler(pmix_status_t codes[], size_t ncodes,
5                             pmix_info_t info[], size_t ninfo,
6                             pmix_notification_fn_t evhdlr,
7                             pmix_hdlr_reg_cbfunc_t cbfunc,
8                             void *cbdata);

```

↕ C ↕

9 *PMIx v2.0* [pmix_event_notification_cbfunc_fn_t](#)

↕ C ↕

```

10  typedef void (*pmix_event_notification_cbfunc_fn_t)
11      (pmix_status_t status,
12       pmix_info_t *results, size_t nresults,
13       pmix_op_cbfunc_t cbfunc, void *thiscbdata,
14       void *notification_cbdata);

```

↕ C ↕

Related Attributes

PMIX_EVENT_HDLR_NAME "pmix.evname" (char*)

String name identifying this handler.

PMIX_EVENT_HDLR_FIRST "pmix.evfirst" (bool)

Invoke this event handler before any other handlers.

PMIX_EVENT_HDLR_LAST "pmix.evlast" (bool)

Invoke this event handler after all other handlers have been called.

PMIX_EVENT_HDLR_FIRST_IN_CATEGORY "pmix.evfirstcat" (bool)

Invoke this event handler before any other handlers in this category.

PMIX_EVENT_HDLR_LAST_IN_CATEGORY "pmix.evlastcat" (bool)

Invoke this event handler after all other handlers in this category have been called.

PMIX_EVENT_HDLR_BEFORE "pmix.evbefore" (char*)

Put this event handler immediately before the one specified in the (char*) value.

PMIX_EVENT_HDLR_AFTER "pmix.evafter" (char*)

Put this event handler immediately after the one specified in the (char*) value.

PMIX_EVENT_HDLR_APPEND "pmix.evappend" (bool)

Append this handler to the precedence list within its category.

Related Constants

PMIX_EVENT_NO_ACTION_TAKEN

PMIX_EVENT_PARTIAL_ACTION_TAKEN

PMIX_EVENT_ACTION_DEFERRED

B.4 MPI Sessions

B.4.1 Use Case Summary

MPI Sessions addresses a number of the limitations of the current MPI programming model.

Among the immediate problems MPI Sessions is intended to address are the following:

- MPI cannot be initialized within an MPI process from different application components without a priori knowledge or coordination,
- MPI cannot be initialized more than once, and MPI cannot be reinitialized after MPI finalize has been called.
- With MPI Sessions, an application no longer needs to explicitly call **MPI_Init** to make use of MPI, but rather can use a Session to only initialize MPI resources for specific communication needs.




Figure B.4.: MPI Communicator from MPI Session Handle using PMIx

- Unless the MPI process explicitly calls `MPI_Init`, there is also no explicit `MPI_COMM_WORLD` communicator. Sessions can be created and destroyed multiple times in an MPI process.

B.4.2 Use Case Details

A PMIx Process Set (PSET) is a user-provided or host environment assigned label associated with a given set of application processes. Processes can belong to multiple process sets at a time. Definition of a PMIx process set typically occurs at time of application execution - e.g., on a command line: `prun -n 4 -pset ocean myoceanapp : -n 3 -pset ice myiceapp`

PMIx PSETs are used for query functions (`MPI_SESSION_GET_NUM_PSETS`, `MPI_SESSION_GET_NTH_PSET`) and to create `MPI_GROUP` from a process set name.

In OpenMPI's MPI Sessions prototype, PMIx groups are used during creation of `MPI_COMM` from an `MPI_GROUP`. The PMIx group constructor returns a 64-bit PMIx Group Context Identifier (PGCID) that is guaranteed to be unique for the duration of an allocation (in the case of a batch managed environment). This PGCID could be used as a direct replacement for the existing unique identifiers for communicators in MPI (E.g. Communicator Identifiers (CIDs) in Open MPI), but may have performance implications.

1 There is an important distinction between process sets and process groups. The process set
2 identifiers are set by the host environment and currently there are no PMIx APIs provided by which
3 an application can change a process set membership. In contrast, PMIx process groups can only be
4 defined dynamically by the application.

5 **Related Interfaces**

6 *PMIx v1.0* **PMIx_Get**

```
▼ _____ C _____ ▼  
7 pmix_status_t  
8 PMIx_Get(const pmix_proc_t *proc, const pmix_key_t key,  
9         const pmix_info_t info[], size_t ninfo,  
10        pmix_value_t **val);  
▲ _____ C _____ ▲
```

11 *PMIx v4.0* **PMIx_Group_construct**

```
▼ _____ C _____ ▼  
12 pmix_status_t  
13 PMIx_Group_construct(const char grp[],  
14                    const pmix_proc_t procs[], size_t nprocs,  
15                    const pmix_info_t directives[],  
16                    size_t ndirs,  
17                    pmix_info_t **results,  
18                    size_t *nresults);  
▲ _____ C _____ ▲
```

19 **Related Attributes**

20 **PMIX_PSET_NAMES** "pmix.pset.nms" (pmix_data_array_t*)

21 Returns an array of **char*** string names of the process sets in which the given process is a
22 member.

23 **PMIX_QUERY_NUM_GROUPS** "pmix.qry.pgrpnum" (size_t)

24 Return the number of process groups defined in the specified range (defaults to session).
25 OPTIONAL QUALIFIERS: **PMIX_RANGE**.

26 **PMIX_QUERY_GROUP_NAMES** "pmix.qry.pgrp" (pmix_data_array_t*)

27 Return a **pmix_data_array_t** containing an array of string names of the process groups
28 defined in the specified range (defaults to session). OPTIONAL QUALIFIERS:
29 **PMIX_RANGE**.

30 **PMIX_QUERY_GROUP_MEMBERSHIP**

31 "pmix.qry.pgrpmems" (pmix_data_array_t*)

32 Return a **pmix_data_array_t** of **pmix_proc_t** containing the members of the
33 specified process group. REQUIRED QUALIFIERS: **PMIX_GROUP_ID**.

Related Constants

[PMIX_SUCCESS](#)

[PMIX_ERR_NOT_SUPPORTED](#)

B.5 Cross-Version Compatibility

B.5.1 Use Case Summary

The PMIx interface serves as a conduit between clients (e.g., MPI libraries), tools (e.g., debuggers), and servers (e.g., RMs). As such, it is probable that a process operating in one of these roles (e.g., as a client or tool) is running a different version of the same PMIx implementation than the process with which it is communicating that is operating in a different role (e.g., as a server). For processes running in containers cross-version compatibility is especially important because the container image and the system software levels will naturally evolve and drift apart. As such, there is a need for PMIx implementations to provide cross-version compatibility.

The responsibility for providing cross-version compatibility is a feature of a specific PMIx implementation and not necessarily of the PMIx standard. The PMIx standard must strive to enable, and never limit, both the cross-version compatibility in any given PMIx implementation, and the ability for a PMIx consumer to adapt to cross-version differences in capabilities.

This use case is focused on cross-version compatibility between different versions of the same PMIx implementation and not between different PMIx implementations.

Cross-version compatibility responsibilities are not restricted to PMIx, but a general issue for any library that coordinates across multiple processes. This includes, but not limited to, client/server libraries, and libraries with a user-space and kernel-space component (e.g., high-performance interconnect libraries).

B.5.2 Use Case Details

There are three scenarios that a PMIx implementation and a PMIx consumer must consider. These scenarios use a PMIx Server and a PMIx Client for clarity, though the scenarios also apply to PMIx Tools.

1. **PMIx Server version matches PMIx Client version:** No cross-version considerations are necessary since they are running the same version.
2. **PMIx Server version is older than PMIx Client version:** The implementation must negotiate capabilities during the initial handshake.
This scenario is common if the (possibly containerized) PMIx client application is being run on an established system that does not update as frequently as the application requires. Thus the PMIx Server in the RM is locked to an older version of that PMIx implementation.

1 3. **PMIx Server version is newer than PMIx Client version:** The implementation must negotiate
2 capabilities during the initial handshake.

3 This scenario is common if the (possibly containerized) PMIx client application is being run
4 after a system software upgrade on the system. Thus the PMIx Server in the RM has been
5 upgraded to a newer version of that PMIx implementation and the client is still linked against the
6 older version.

7 When the two PMIx-enabled processes first connect to each other they need to first check the
8 version of the library that they are each running. This handshake often occurs during initialization
9 (though it could occur on a per-operation basis depending on the specific PMIx implementation),
10 for example during the following operations:

- 11 • PMIx Clients: [PMIx_Init](#)
- 12 • PMIx Tools: [PMIx_tool_init](#), [PMIx_tool_attach_to_server](#)
- 13 • PMIx Servers: [PMIx_server_init](#), [pmix_server_client_connected2_fn_t](#),
14 [pmix_server_tool_connection_fn_t](#)

15 Commonly this cross-version handshake occurs completely transparently to the consumers of the
16 PMIx interface since it happens inside a specific PMIx implementation of these interfaces.
17 However, during the negotiation, some features available in one version might not be available in
18 the other. The consumer of the PMIx interface should always be prepared to receive the
19 [PMIX_ERR_NOT_SUPPORTED](#) error code from a PMIx interface call that the other side either
20 does not support or is not available in the version of the library with which they are linked. After
21 connecting to another PMIx entity, the consumer of the PMIx interface can use the
22 [PMIx_Query_info](#) API to determine supported functionality and adapt accordingly.

23 Related Interfaces

24 *PMIx v1.2*

[PMIx_Init](#)

```
25 pmix_status_t  
26 PMIx_Init(pmix_proc_t *proc,  
27           pmix_info_t info[], size_t ninfo)
```

28 *PMIx v2.0*

[PMIx_tool_init](#)

```
29 pmix_status_t  
30 PMIx_tool_init(pmix_proc_t *proc,  
31               pmix_info_t info[], size_t ninfo);
```

32 *PMIx v4.0*

[PMIx_tool_attach_to_server](#)

```

1      pmix_status_t
2      PMIx_tool_attach_to_server(pmix_proc_t *proc,
3                                  pmix_proc_t *server,
4                                  pmix_info_t info[],
5                                  size_t ninfo);
6 PMIx v1.0 PMIx_server_init
7
8      pmix_status_t
9      PMIx_server_init(pmix_server_module_t *module,
10                       pmix_info_t info[], size_t ninfo);
11 PMIx v4.0 pmix_server_client_connected2_fn_t
12
13      typedef pmix_status_t (*pmix_server_client_connected2_fn_t)(
14                          const pmix_proc_t *proc,
15                          void* server_object,
16                          pmix_info_t info[], size_t ninfo,
17                          pmix_op_cbfunc_t cbfunc,
18                          void *cbdata)
19
20 PMIx v2.0 pmix_server_tool_connection_fn_t
21
22      typedef void (*pmix_server_tool_connection_fn_t)(
23                          pmix_info_t info[], size_t ninfo,
24                          pmix_tool_connection_cbfunc_t cbfunc,
25                          void *cbdata);
26 PMIx v4.0 PMIx_Query_info
27
28      pmix_status_t
29      PMIx_Query_info(pmix_query_t queries[], size_t nqueries,
30                     pmix_info_t *info[], size_t *ninfo);

```

1
2
3

Related Constants

`PMIX_SUCCESS`

`PMIX_ERR_NOT_SUPPORTED`

Unofficial Draft

APPENDIX C

Revision History

C.1 Version 1.0: June 12, 2015

The PMIx version 1.0 *ad hoc* standard was defined in a set of header files as part of the v1.0.0 release of the OpenPMIx library prior to the creation of the formal PMIx 2.0 standard. Below are a summary listing of the interfaces defined in the 1.0 headers.

- Client APIs
 - `PMIx_Init`, `PMIx_Initialized`, `PMIx_Abort`, `PMIx_Finalize`
 - `PMIx_Put`, `PMIx_Commit`,
 - `PMIx_Fence`, `PMIx_Fence_nb`
 - `PMIx_Get`, `PMIx_Get_nb`
 - `PMIx_Publish`, `PMIx_Publish_nb`
 - `PMIx_Lookup`, `PMIx_Lookup_nb`
 - `PMIx_Unpublish`, `PMIx_Unpublish_nb`
 - `PMIx_Spawn`, `PMIx_Spawn_nb`
 - `PMIx_Connect`, `PMIx_Connect_nb`
 - `PMIx_Disconnect`, `PMIx_Disconnect_nb`
 - `PMIx_Resolve_nodes`, `PMIx_Resolve_peers`
- Server APIs
 - `PMIx_server_init`, `PMIx_server_finalize`
 - `PMIx_generate_regex`, `PMIx_generate_ppn`
 - `PMIx_server_register_nspace`, `PMIx_server_deregister_nspace`
 - `PMIx_server_register_client`, `PMIx_server_deregister_client`
 - `PMIx_server_setup_fork`, `PMIx_server_dmodex_request`
- Common APIs
 - `PMIx_Get_version`, `PMIx_Store_internal`, `PMIx_Error_string`
 - `PMIx_Register_errhandler`, `PMIx_Deregister_errhandler`, `PMIx_Notify_error`

The `PMIx_Init` API was subsequently modified in the v1.1.0 release of that library.

1 C.2 Version 2.0: Sept. 2018

2 The following APIs were introduced in v2.0 of the PMIx Standard:

- 3 • Client APIs
 - 4 – `PMIx_Query_info_nb`, `PMIx_Log_nb`
 - 5 – `PMIx_Allocation_request_nb`, `PMIx_Job_control_nb`,
 - 6 `PMIx_Process_monitor_nb`, `PMIx_Heartbeat`
- 7 • Server APIs
 - 8 – `PMIx_server_setup_application`, `PMIx_server_setup_local_support`
- 9 • Tool APIs
 - 10 – `PMIx_tool_init`, `PMIx_tool_finalize`
- 11 • Common APIs
 - 12 – `PMIx_Register_event_handler`, `PMIx_Deregister_event_handler`
 - 13 – `PMIx_Notify_event`
 - 14 – `PMIx_Proc_state_string`, `PMIx_Scope_string`
 - 15 – `PMIx_Persistence_string`, `PMIx_Data_range_string`
 - 16 – `PMIx_Info_directives_string`, `PMIx_Data_type_string`
 - 17 – `PMIx_Alloc_directive_string`
 - 18 – `PMIx_Data_pack`, `PMIx_Data_unpack`, `PMIx_Data_copy`
 - 19 – `PMIx_Data_print`, `PMIx_Data_copy_payload`

20 C.2.1 Removed/Modified APIs

21 The `PMIx_Init` API was modified in v2.0 of the standard from its *ad hoc* v1.0 signature to
22 include passing of a `pmix_info_t` array for flexibility and “future-proofing” of the API. In
23 addition, the `PMIx_Notify_error`, `PMIx_Register_errhandler`, and
24 `PMIx_Deregister_errhandler` APIs were replaced. This pre-dated official adoption of
25 PMIx as a Standard.

26 C.2.2 Deprecated constants

27 The following constants were deprecated in v2.0:

- 28 `PMIX_MODEX`
- 29 `PMIX_INFO_ARRAY`

1 C.2.3 Deprecated attributes

2 The following attributes were deprecated in v2.0:

- 3 **PMIX_ERROR_NAME** "pmix.errname" (pmix_status_t)
4 Specific error to be notified
- 5 **PMIX_ERROR_GROUP_COMM** "pmix.errgroup.comm" (bool)
6 Set true to get comm errors notification
- 7 **PMIX_ERROR_GROUP_ABORT** "pmix.errgroup.abort" (bool)
8 Set true to get abort errors notification
- 9 **PMIX_ERROR_GROUP_MIGRATE** "pmix.errgroup.migrate" (bool)
10 Set true to get migrate errors notification
- 11 **PMIX_ERROR_GROUP_RESOURCE** "pmix.errgroup.resource" (bool)
12 Set true to get resource errors notification
- 13 **PMIX_ERROR_GROUP_SPAWN** "pmix.errgroup.spawn" (bool)
14 Set true to get spawn errors notification
- 15 **PMIX_ERROR_GROUP_NODE** "pmix.errgroup.node" (bool)
16 Set true to get node status notification
- 17 **PMIX_ERROR_GROUP_LOCAL** "pmix.errgroup.local" (bool)
18 Set true to get local errors notification
- 19 **PMIX_ERROR_GROUP_GENERAL** "pmix.errgroup.gen" (bool)
20 Set true to get notified of generic errors
- 21 **PMIX_ERROR_HANDLER_ID** "pmix.errhandler.id" (int)
22 Errhandler reference id of notification being reported

23 C.3 Version 2.1: Dec. 2018

24 The v2.1 update includes clarifications and corrections from the v2.0 document, plus addition of
25 examples:

- 26 • Clarify description of [PMIx_Connect](#) and [PMIx_Disconnect](#) APIs.
- 27 • Explain that values for the [PMIX_COLLECTIVE_ALGO](#) are environment-dependent
- 28 • Identify the namespace/rank values required for retrieving attribute-associated information using
29 the [PMIx_Get](#) API
- 30 • Provide definitions for *session*, *job*, *application*, and other terms used throughout the document
- 31 • Clarify definitions of [PMIX_UNIV_SIZE](#) versus [PMIX_JOB_SIZE](#)
- 32 • Clarify server module function return values
- 33 • Provide examples of the use of [PMIx_Get](#) for retrieval of information
- 34 • Clarify the use of [PMIx_Get](#) versus [PMIx_Query_info_nb](#)
- 35 • Clarify return values for non-blocking APIs and emphasize that callback functions must not be
36 invoked prior to return from the API
- 37 • Provide detailed example for construction of the [PMIx_server_register_namespace](#) input
38 information array

- 1 • Define information levels (e.g., *session* vs *job*) and associated attributes for both storing and
- 2 retrieving values
- 3 • Clarify roles of PMIx server library and host environment for collective operations
- 4 • Clarify definition of `PMIX_UNIV_SIZE`

5 C.4 Version 2.2: Jan 2019

6 The v2.2 update includes the following clarifications and corrections from the v2.1 document:

- 7 • Direct modex upcall function (`pmix_server_dmodex_req_fn_t`) cannot complete
- 8 atomically as the API cannot return the requested information except via the provided callback
- 9 function
- 10 • Add missing `pmix_data_array_t` definition and support macros
- 11 • Add a rule divider between implementer and host environment required attributes for clarity
- 12 • Add `PMIX_QUERY_QUALIFIERS_CREATE` macro to simplify creation of `pmix_query_t`
- 13 qualifiers
- 14 • Add `PMIX_APP_INFO_CREATE` macro to simplify creation of `pmix_app_t` directives
- 15 • Add flag and `PMIX_INFO_IS_END` macro for marking and detecting the end of a
- 16 `pmix_info_t` array
- 17 • Clarify the allowed hierarchical nesting of the `PMIX_SESSION_INFO_ARRAY`,
- 18 `PMIX_JOB_INFO_ARRAY`, and associated attributes

19 C.5 Version 3.0: Dec. 2018

20 The following APIs were introduced in v3.0 of the PMIx Standard:

- 21 • Client APIs
 - 22 – `PMIx_Log`, `PMIx_Job_control`
 - 23 – `PMIx_Allocation_request`, `PMIx_Process_monitor`
 - 24 – `PMIx_Get_credential`, `PMIx_Validate_credential`
- 25 • Server APIs
 - 26 – `PMIx_server_IOF_deliver`
 - 27 – `PMIx_server_collect_inventory`, `PMIx_server_deliver_inventory`
- 28 • Tool APIs
 - 29 – `PMIx_IOF_pull`, `PMIx_IOF_push`, `PMIx_IOF_deregister`
 - 30 – `PMIx_tool_connect_to_server`
- 31 • Common APIs
 - 32 – `PMIx_IOF_channel_string`

1 The document added a chapter on security credentials, a new section for IO forwarding to the
2 Process Management chapter, and a few blocking forms of previously-existing non-blocking APIs.
3 Attributes supporting the new APIs were introduced, as well as additional attributes for a few
4 existing functions.

5 C.5.1 Removed constants

6 The following constants were removed in v3.0:

7 `PMIX_MODEX`
8 `PMIX_INFO_ARRAY`

9 C.5.2 Deprecated attributes

10 The following attributes were deprecated in v3.0:

11 `PMIX_COLLECTIVE_ALGO_REQD` "pmix.calreqd" (bool)
12 If `true`, indicates that the requested choice of algorithm is mandatory.

13 C.5.3 Removed attributes

14 The following attributes were removed in v3.0:

15 `PMIX_ERROR_NAME` "pmix.errname" (pmix_status_t)
16 Specific error to be notified
17 `PMIX_ERROR_GROUP_COMM` "pmix.errgroup.comm" (bool)
18 Set true to get comm errors notification
19 `PMIX_ERROR_GROUP_ABORT` "pmix.errgroup.abort" (bool)
20 Set true to get abort errors notification
21 `PMIX_ERROR_GROUP_MIGRATE` "pmix.errgroup.migrate" (bool)
22 Set true to get migrate errors notification
23 `PMIX_ERROR_GROUP_RESOURCE` "pmix.errgroup.resource" (bool)
24 Set true to get resource errors notification
25 `PMIX_ERROR_GROUP_SPAWN` "pmix.errgroup.spawn" (bool)
26 Set true to get spawn errors notification
27 `PMIX_ERROR_GROUP_NODE` "pmix.errgroup.node" (bool)
28 Set true to get node status notification
29 `PMIX_ERROR_GROUP_LOCAL` "pmix.errgroup.local" (bool)
30 Set true to get local errors notification
31 `PMIX_ERROR_GROUP_GENERAL` "pmix.errgroup.gen" (bool)
32 Set true to get notified of generic errors
33 `PMIX_ERROR_HANDLER_ID` "pmix.errhandler.id" (int)
34 Errhandler reference id of notification being reported

1 C.6 Version 3.1: Jan. 2019

2 The v3.1 update includes clarifications and corrections from the v3.0 document:

- 3 • Direct modex upcall function (`pmix_server_dmodex_req_fn_t`) cannot complete
- 4 atomically as the API cannot return the requested information except via the provided callback
- 5 function
- 6 • Fix typo in name of `PMIX_FWD_STDDIAG` attribute
- 7 • Correctly identify the information retrieval and storage attributes as “new” to v3 of the standard
- 8 • Add missing `pmix_data_array_t` definition and support macros
- 9 • Add a rule divider between implementer and host environment required attributes for clarity
- 10 • Add `PMIX_QUERY_QUALIFIERS_CREATE` macro to simplify creation of `pmix_query_t`
- 11 `qualifiers`
- 12 • Add `PMIX_APP_INFO_CREATE` macro to simplify creation of `pmix_app_t` directives
- 13 • Add new attributes to specify the level of information being requested where ambiguity may exist
- 14 (see 6.1)
- 15 • Add new attributes to assemble information by its level for storage where ambiguity may exist
- 16 (see 17.2.3.1)
- 17 • Add flag and `PMIX_INFO_IS_END` macro for marking and detecting the end of a
- 18 `pmix_info_t` array
- 19 • Clarify that `PMIX_NUM_SLOTS` is duplicative of (a) `PMIX_UNIV_SIZE` when used at the
- 20 *session* level and (b) `PMIX_MAX_PROCS` when used at the *job* and *application* levels, but leave
- 21 it in for backward compatibility.
- 22 • Clarify difference between `PMIX_JOB_SIZE` and `PMIX_MAX_PROCS`
- 23 • Clarify that `PMIx_server_setup_application` must be called per-*job* instead of
- 24 per-*application* as the name implies. Unfortunately, this is a historical artifact. Note that both
- 25 `PMIX_NODE_MAP` and `PMIX_PROC_MAP` must be included as input in the *info* array provided
- 26 to that function. Further descriptive explanation of the “instant on” procedure will be provided in
- 27 the next version of the PMIx Standard.
- 28 • Clarify how the PMIx server expects data passed to the host by
- 29 `pmix_server_fence_fn_t` should be aggregated across nodes, and provide a code
- 30 snippet example

31 C.7 Version 3.2: Oct. 2020

32 The v3.2 update includes clarifications and corrections from the v3.1 document:

- 33 • Correct an error in the `PMIx_Allocation_request` function signature, and clarify the
- 34 allocation ID attributes
- 35 • Rename the `PMIX_ALLOC_ID` attribute to `PMIX_ALLOC_REQ_ID` to clarify that this is a
- 36 string the user provides as a means to identify their request to query status

- Add a new `PMIX_ALLOC_ID` attribute that contains the identifier (provided by the host environment) for the resulting allocation which can later be used to reference the allocated resources in, for example, a call to `PMIx_Spawn`
- Update the `PMIx_generate_regex` and `PMIx_generate_ppn` descriptions to clarify that the output from these generator functions may not be a NULL-terminated string, but instead could be a byte array of arbitrary binary content.
- Add a new `PMIX_REGEX` constant that represents a regular expression data type.

C.7.1 Deprecated constants

The following constants were deprecated in v3.2:

<code>PMIX_ERR_DATA_VALUE_NOT_FOUND</code>	Data value not found
<code>PMIX_ERR_HANDSHAKE_FAILED</code>	Connection handshake failed
<code>PMIX_ERR_IN_ERRNO</code>	Error defined in <code>errno</code>
<code>PMIX_ERR_INVALID_ARG</code>	Invalid argument
<code>PMIX_ERR_INVALID_ARGS</code>	Invalid arguments
<code>PMIX_ERR_INVALID_KEY</code>	Invalid key
<code>PMIX_ERR_INVALID_KEY_LENGTH</code>	Invalid key length
<code>PMIX_ERR_INVALID_KEYVALP</code>	Invalid key/value pair
<code>PMIX_ERR_INVALID_LENGTH</code>	Invalid argument length
<code>PMIX_ERR_INVALID_NAMESPACE</code>	Invalid namespace
<code>PMIX_ERR_INVALID_NUM_ARGS</code>	Invalid number of arguments
<code>PMIX_ERR_INVALID_NUM_PARSED</code>	Invalid number parsed
<code>PMIX_ERR_INVALID_SIZE</code>	Invalid size
<code>PMIX_ERR_INVALID_VAL</code>	Invalid value
<code>PMIX_ERR_INVALID_VAL_LENGTH</code>	Invalid value length
<code>PMIX_ERR_NOT_IMPLEMENTED</code>	Not implemented
<code>PMIX_ERR_PACK_MISMATCH</code>	Pack mismatch
<code>PMIX_ERR_PROC_ENTRY_NOT_FOUND</code>	Process not found
<code>PMIX_ERR_PROC_REQUESTED_ABORT</code>	Process is already requested to abort
<code>PMIX_ERR_READY_FOR_HANDSHAKE</code>	Ready for handshake
<code>PMIX_ERR_SERVER_FAILED_REQUEST</code>	Failed to connect to the server
<code>PMIX_ERR_SERVER_NOT_AVAIL</code>	Server is not available
<code>PMIX_ERR_SILENT</code>	Silent error
<code>PMIX_GDS_ACTION_COMPLETE</code>	The Global Data Storage (GDS) action has completed
<code>PMIX_NOTIFY_ALLOC_COMPLETE</code>	Notify that a requested allocation operation is complete
	- the result of the request will be included in the <i>info</i> array

1 C.7.2 Deprecated attributes

2 The following attributes were deprecated in v3.2:

3 **PMIX_ARCH** "pmix.arch" (uint32_t)

4 Architecture flag.

5 **PMIX_COLLECTIVE_ALGO** "pmix.calgo" (char*)

6 Comma-delimited list of algorithms to use for the collective operation. PMIx does not
7 impose any requirements on a host environment's collective algorithms. Thus, the
8 acceptable values for this attribute will be environment-dependent - users are encouraged to
9 check their host environment for supported values.

10 **PMIX_DSTPATH** "pmix.dstpath" (char*)

11 Path to shared memory data storage (dstore) files. Deprecated from Standard as being
12 implementation specific.

13 **PMIX_HWLOC_HOLE_KIND** "pmix.hwlocholek" (char*)

14 Kind of VM "hole" HWLOC should use for shared memory

15 **PMIX_HWLOC_SHARE_TOPO** "pmix.hwlocsh" (bool)

16 Share the HWLOC topology via shared memory

17 **PMIX_HWLOC_SHMEM_ADDR** "pmix.hwlocaddr" (size_t)

18 Address of the HWLOC shared memory segment.

19 **PMIX_HWLOC_SHMEM_FILE** "pmix.hwlocfile" (char*)

20 Path to the HWLOC shared memory file.

21 **PMIX_HWLOC_SHMEM_SIZE** "pmix.hwlocsize" (size_t)

22 Size of the HWLOC shared memory segment.

23 **PMIX_HWLOC_XML_V1** "pmix.hwlocxml1" (char*)

24 XML representation of local topology using HWLOC's v1.x format.

25 **PMIX_HWLOC_XML_V2** "pmix.hwlocxml2" (char*)

26 XML representation of local topology using HWLOC's v2.x format.

27 **PMIX_LOCAL_TOPO** "pmix.ltopo" (char*)

28 XML representation of local node topology.

29 **PMIX_MAPPER** "pmix.mapper" (char*)

30 Mapping mechanism to use for placing spawned processes - when accessed using
31 [PMIx_Get](#), use the [PMIX_RANK_WILDCARD](#) value for the rank to discover the mapping
32 mechanism used for the provided namespace.

33 **PMIX_MAP_BLOB** "pmix.mblob" (pmix_byte_object_t)

34 Packed blob of process location.

35 **PMIX_NON_PMI** "pmix.nonpmi" (bool)

36 Spawned processes will not call [PMIx_Init](#).

37 **PMIX_PROC_BLOB** "pmix.pblob" (pmix_byte_object_t)

38 Packed blob of process data.

39 **PMIX_PROC_URI** "pmix.puri" (char*)

40 URI containing contact information for the specified process.

41 **PMIX_TOPOLOGY_FILE** "pmix.topo.file" (char*)

42 Full path to file containing XML topology description

1 **PMIX_TOPOLOGY_SIGNATURE** "pmix.toposig" (char*)
2 Topology signature string.
3 **PMIX_TOPOLOGY_XML** "pmix.topo.xml" (char*)
4 XML-based description of topology

5 C.8 Version 4.0: Dec. 2020

6 NOTE: The PMIx Standard document has undergone significant reorganization in an effort to
7 become more user-friendly. Highlights include:

- 8 • Moving all added, deprecated, and removed items to this revision log section to make them more
9 visible
- 10 • Co-locating constants and attribute definitions with the primary API that uses them - citations
11 and hyperlinks are retained elsewhere
- 12 • Splitting the Key-Value Management chapter into separate chapters on the use of reserved keys,
13 non-reserved keys, and non-process-related key-value data exchange
- 14 • Creating a new chapter on synchronization and data access methods
- 15 • Removing references to specific implementations of PMIx and to implementation-specific
16 features and/or behaviors

17 In addition to the reorganization, the following changes were introduced in v4.0 of the PMIx
18 Standard:

- 19 • Clarified that the **PMIx_Fence_nb** operation can immediately return
20 **PMIX_OPERATION_SUCCEEDED** in lieu of passing the request to a PMIx server if only the
21 calling process is involved in the operation
- 22 • Added the **PMIx_Register_attributes** API by which a host environment can register the
23 attributes it supports for each server-to-host operation
- 24 • Added the ability to query supported attributes from the PMIx tool, client and server libraries, as
25 well as the host environment via the new **pmix_regattr_t** structure. Both human-readable
26 and machine-parsable output is supported. New attributes to support this operation include:
 - 27 – **PMIX_CLIENT_ATTRIBUTES**, **PMIX_SERVER_ATTRIBUTES**,
 - 28 **PMIX_TOOL_ATTRIBUTES**, and **PMIX_HOST_ATTRIBUTES** to identify which library
29 supports the attribute; and
 - 30 – **PMIX_MAX_VALUE**, **PMIX_MIN_VALUE**, and **PMIX_ENUM_VALUE** to provide
31 machine-parsable description of accepted values
- 32 • Add **PMIX_APP_WILDCARD** to reference all applications within a given job
- 33 • Fix signature of blocking APIs **PMIx_Allocation_request**, **PMIx_Job_control**,
34 **PMIx_Process_monitor**, **PMIx_Get_credential**, and
35 **PMIx_Validate_credential** to allow return of results
- 36 • Update description to provide an option for blocking behavior of the
37 **PMIx_Register_event_handler**, **PMIx_Deregister_event_handler**,
38 **PMIx_Notify_event**, **PMIx_IOF_pull**, **PMIx_IOF_deregister**, and
39 **PMIx_IOF_push** APIs. The need for blocking forms of these functions was not initially

1 anticipated but has emerged over time. For these functions, the return value is sufficient to
2 provide the caller with information otherwise returned via callback. Thus, use of a **NULL** value
3 as the callback function parameter was deemed a minimal disruption method for providing the
4 desired capability

- 5 • Added a chapter on fabric support that includes new APIs, datatypes, and attributes
- 6 • Added a chapter on process sets and groups that includes new APIs and attributes
- 7 • Added APIs and a new datatypes to support generation and parsing of PMIx locality and cpuset
8 strings
- 9 • Added a new chapter on tools that provides deeper explanation on their operation and collecting
10 all tool-relevant definitions into one location. Also introduced two new APIs and removed
11 restriction that limited tools to being connected to only one server at a time.
- 12 • Extended behavior of `PMIx_server_init` to scalably expose the topology description to the
13 local clients. This includes creating any required shared memory backing stores and/or XML
14 representations, plus ensuring that all necessary key-value pairs for clients to access the
15 description are included in the job-level information provided to each client.
- 16 • Added a new API by which the host can manually progress the PMIx library in lieu of the
17 library's own progress thread. s

18 The above changes included introduction of the following APIs and data types:

19 • Client APIs

- 20 – `PMIx_Group_construct`, `PMIx_Group_construct_nb`
- 21 – `PMIx_Group_destruct`, `PMIx_Group_destruct_nb`
- 22 – `PMIx_Group_invite`, `PMIx_Group_invite_nb`
- 23 – `PMIx_Group_join`, `PMIx_Group_join_nb`
- 24 – `PMIx_Group_leave`, `PMIx_Group_leave_nb`
- 25 – `PMIx_Get_relative_locality`, `PMIx_Load_topology`
- 26 – `PMIx_Parse_cpuset_string`, `PMIx_Get_cpuset`
- 27 – `PMIx_Link_state_string`, `PMIx_Job_state_string`
- 28 – `PMIx_Device_type_string`
- 29 – `PMIx_Fabric_register`, `PMIx_Fabric_register_nb`
- 30 – `PMIx_Fabric_update`, `PMIx_Fabric_update_nb`
- 31 – `PMIx_Fabric_deregister`, `PMIx_Fabric_deregister_nb`
- 32 – `PMIx_Compute_distances`, `PMIx_Compute_distances_nb`
- 33 – `PMIx_Get_attribute_string`, `PMIx_Get_attribute_name`
- 34 – `PMIx_Progress`

35 • Server APIs

- 36 – `PMIx_server_generate_locality_string`
- 37 – `PMIx_Register_attributes`
- 38 – `PMIx_server_define_process_set`, `PMIx_server_delete_process_set`
- 39 – `pmix_server_grp_fn_t`, `pmix_server_fabric_fn_t`
- 40 – `pmix_server_client_connected2_fn_t`

- 1 - `PMIx_server_generate_cpuset_string`
- 2 - `PMIx_server_register_resources, PMIx_server_deregister_resources`
- 3 • Tool APIs
- 4 - `PMIx_tool_disconnect`
- 5 - `PMIx_tool_set_server`
- 6 - `PMIx_tool_attach_to_server`
- 7 - `PMIx_tool_get_servers`
- 8 • Data types
- 9 - `pmix_regattr_t`
- 10 - `pmix_cpuset_t`
- 11 - `pmix_topology_t`
- 12 - `pmix_locality_t`
- 13 - `pmix_bind_envelope_t`
- 14 - `pmix_group_opt_t`
- 15 - `pmix_group_operation_t`
- 16 - `pmix_fabric_t`
- 17 - `pmix_device_distance_t`
- 18 - `pmix_coord_t`
- 19 - `pmix_coord_view_t`
- 20 - `pmix_geometry_t`
- 21 - `pmix_link_state_t`
- 22 - `pmix_job_state_t`
- 23 - `pmix_device_type_t`
- 24 • Callback functions
- 25 - `pmix_device_dist_cbfunc_t`

26 **C.8.1 Added Constants**

27 **General error constants**

- 28 `PMIX_ERR_EXISTS_OUTSIDE_SCOPE`
- 29 `PMIX_ERR_PARAM_VALUE_NOT_SUPPORTED`
- 30 `PMIX_ERR_EMPTY`
- 31

Data type constants

PMIX_COORD
PMIX_REGATTR
PMIX_REGEX
PMIX_JOB_STATE
PMIX_LINK_STATE
PMIX_PROC_CPuset
PMIX_GEOMETRY
PMIX_DEVICE_DIST
PMIX_ENDPOINT
PMIX_TOPO
PMIX_DEVTYPE
PMIX_LOCTYPE
PMIX_DATA_TYPE_MAX
PMIX_COMPRESSED_BYTE_OBJECT

Info directives

PMIX_INFO_REQD_PROCESSED

Server constants

PMIX_ERR_REPEAT_ATTR_REGISTRATION

Job-Mgmt constants

PMIX_ERR_CONFLICTING_CLEANUP_DIRECTIVES

Publish constants

PMIX_ERR_DUPLICATE_KEY

Tool constants

PMIX_LAUNCHER_READY
PMIX_ERR_IOF_FAILURE
PMIX_ERR_IOF_COMPLETE
PMIX_EVENT_JOB_START
PMIX_LAUNCH_COMPLETE
PMIX_EVENT_JOB_END
PMIX_EVENT_SESSION_START
PMIX_EVENT_SESSION_END
PMIX_ERR_PROC_TERM_WO_SYNC
PMIX_ERR_JOB_CANCELED
PMIX_ERR_JOB_ABORTED

1 PMIX_ERR_JOB_KILLED_BY_CMD
2 PMIX_ERR_JOB_ABORTED_BY_SIG
3 PMIX_ERR_JOB_TERM_WO_SYNC
4 PMIX_ERR_JOB_SENSOR_BOUND_EXCEEDED
5 PMIX_ERR_JOB_NON_ZERO_TERM
6 PMIX_ERR_JOB_ABORTED_BY_SYS_EVENT
7 PMIX_DEBUG_WAITING_FOR_NOTIFY
8 PMIX_DEBUGGER_RELEASE
9

10 **Fabric constants**

11 PMIX_FABRIC_UPDATE_PENDING
12 PMIX_FABRIC_UPDATED
13 PMIX_FABRIC_UPDATE_ENDPOINTS
14 PMIX_COORD_VIEW_UNDEF
15 PMIX_COORD_LOGICAL_VIEW
16 PMIX_COORD_PHYSICAL_VIEW
17 PMIX_LINK_STATE_UNKNOWN
18 PMIX_LINK_DOWN
19 PMIX_LINK_UP
20 PMIX_FABRIC_REQUEST_INFO
21 PMIX_FABRIC_UPDATE_INFO
22

23 **Sets-Groups constants**

24 PMIX_PROCESS_SET_DEFINE
25 PMIX_PROCESS_SET_DELETE
26 PMIX_GROUP_INVITED
27 PMIX_GROUP_LEFT
28 PMIX_GROUP_MEMBER_FAILED
29 PMIX_GROUP_INVITE_ACCEPTED
30 PMIX_GROUP_INVITE_DECLINED
31 PMIX_GROUP_INVITE_FAILED
32 PMIX_GROUP_MEMBERSHIP_UPDATE
33 PMIX_GROUP_CONSTRUCT_ABORT
34 PMIX_GROUP_CONSTRUCT_COMPLETE
35 PMIX_GROUP_LEADER_FAILED
36 PMIX_GROUP_LEADER_SELECTED
37 PMIX_GROUP_CONTEXT_ID_ASSIGNED
38

39 **Process-Mgmt constants**

40 PMIX_ERR_JOB_ALLOC_FAILED
41 PMIX_ERR_JOB_APP_NOT_EXECUTABLE

1 `PMIX_ERR_JOB_NO_EXE_SPECIFIED`
2 `PMIX_ERR_JOB_FAILED_TO_MAP`
3 `PMIX_ERR_JOB_FAILED_TO_LAUNCH`
4 `PMIX_LOCALITY_UNKNOWN`
5 `PMIX_LOCALITY_NONLOCAL`
6 `PMIX_LOCALITY_SHARE_HWTHREAD`
7 `PMIX_LOCALITY_SHARE_CORE`
8 `PMIX_LOCALITY_SHARE_L1CACHE`
9 `PMIX_LOCALITY_SHARE_L2CACHE`
10 `PMIX_LOCALITY_SHARE_L3CACHE`
11 `PMIX_LOCALITY_SHARE_PACKAGE`
12 `PMIX_LOCALITY_SHARE_NUMA`
13 `PMIX_LOCALITY_SHARE_NODE`
14

15 Events

16 `PMIX_EVENT_SYS_BASE`
17 `PMIX_EVENT_NODE_DOWN`
18 `PMIX_EVENT_NODE_OFFLINE`
19 `PMIX_EVENT_SYS_OTHER`
20

21 C.8.2 Added Attributes

22 Sync-Access attributes

23 `PMIX_COLLECT_GENERATED_JOB_INFO` "pmix.collect.gen" (bool)

24 Collect all job-level information (i.e., reserved keys) that was locally generated by PMIx
25 servers. Some job-level information (e.g., distance between processes and fabric devices) is
26 best determined on a distributed basis as it primarily pertains to local processes. Should
27 remote processes need to access the information, it can either be obtained collectively using
28 the `PMIx_Fence` operation with this directive, or can be retrieved one peer at a time using
29 `PMIx_Get` without first having performed the job-wide collection.

30 `PMIX_ALL_CLONES_PARTICIPATE` "pmix.clone.part" (bool)

31 All *clones* of the calling process must participate in the collective operation.

32 `PMIX_GET_POINTER_VALUES` "pmix.get.pntrs" (bool)

33 Request that any pointers in the returned value point directly to values in the key-value store.
34 The user *must not* release any returned data pointers.

35 `PMIX_GET_STATIC_VALUES` "pmix.get.static" (bool)

36 Request that the data be returned in the provided storage location. The caller is responsible
37 for destructing the `pmix_value_t` using the `PMIX_VALUE_DESTRUCT` macro when
38 done.

39 `PMIX_GET_REFRESH_CACHE` "pmix.get.refresh" (bool)

1 When retrieving data for a remote process, refresh the existing local data cache for the
2 process in case new values have been put and committed by the process since the last refresh.
3 Local process information is assumed to be automatically updated upon posting by the
4 process. A **NULL** key will cause all values associated with the process to be refreshed -
5 otherwise, only the indicated key will be updated. A process rank of
6 **PMIX_RANK_WILDCARD** can be used to update job-related information in dynamic
7 environments. The user is responsible for subsequently updating refreshed values they may
8 have cached in their own local memory.

9 **PMIX_QUERY_RESULTS** "pmix.qry.res" (pmix_data_array_t)

10 Contains an array of query results for a given **pmix_query_t** passed to the
11 **PMIx_Query_info** APIs. If qualifiers were included in the query, then the first element
12 of the array shall be the **PMIX_QUERY_QUALIFIERS** attribute containing those qualifiers.
13 Each of the remaining elements of the array is a **pmix_info_t** containing the query key
14 and the corresponding value returned by the query. This attribute is solely for reporting
15 purposes and cannot be used in **PMIx_Get** or other query operations.

16 **PMIX_QUERY_QUALIFIERS** "pmix.qry.qual" (pmix_data_array_t)

17 Contains an array of qualifiers that were included in the query that produced the provided
18 results. This attribute is solely for reporting purposes and cannot be used in **PMIx_Get** or
19 other query operations.

20 **PMIX_QUERY_SUPPORTED_KEYS** "pmix.qry.keys" (char*)

21 Returns comma-delimited list of keys supported by the query function. NO QUALIFIERS.

22 **PMIX_QUERY_SUPPORTED_QUALIFIERS** "pmix.qry.qual" (char*)

23 Return comma-delimited list of qualifiers supported by a query on the provided key, instead
24 of actually performing the query on the key. NO QUALIFIERS.

25 **PMIX_QUERY_NAMESPACE_INFO** "pmix.qry.nsinfo" (pmix_data_array_t*)

26 Return an array of active namespace information - each element will itself contain an array
27 including the namespace plus the command line of the application executing within it.
28 OPTIONAL QUALIFIERS: **PMIX_NAMESPACE** of specific namespace whose info is being
29 requested.

30 **PMIX_QUERY_ATTRIBUTE_SUPPORT** "pmix.qry.attrs" (bool)

31 Query list of supported attributes for specified APIs. REQUIRED QUALIFIERS: one or
32 more of **PMIX_CLIENT_FUNCTIONS**, **PMIX_SERVER_FUNCTIONS**,
33 **PMIX_TOOL_FUNCTIONS**, and **PMIX_HOST_FUNCTIONS**.

34 **PMIX_QUERY_AVAIL_SERVERS** "pmix.qry.asrvrs" (pmix_data_array_t*)

35 Return an array of **pmix_info_t**, each element itself containing a
36 **PMIX_SERVER_INFO_ARRAY** entry holding all available data for a server on this node to
37 which the caller might be able to connect.

38 **PMIX_SERVER_INFO_ARRAY** "pmix.srv.arr" (pmix_data_array_t)

1 Array of `pmix_info_t` about a given server, starting with its `PMIX_NAMESPACE` and
2 including at least one of the rendezvous-required pieces of information.

3 **PMIX_CLIENT_FUNCTIONS** "`pmix.client.fns`" (`bool`)

4 Request a list of functions supported by the PMIx client library.

5 **PMIX_CLIENT_ATTRIBUTES** "`pmix.client.attrs`" (`bool`)

6 Request attributes supported by the PMIx client library.

7 **PMIX_SERVER_FUNCTIONS** "`pmix.srvr.fns`" (`bool`)

8 Request a list of functions supported by the PMIx server library.

9 **PMIX_SERVER_ATTRIBUTES** "`pmix.srvr.attrs`" (`bool`)

10 Request attributes supported by the PMIx server library.

11 **PMIX_HOST_FUNCTIONS** "`pmix.srvr.fns`" (`bool`)

12 Request a list of functions supported by the host environment.

13 **PMIX_HOST_ATTRIBUTES** "`pmix.host.attrs`" (`bool`)

14 Request attributes supported by the host environment.

15 **PMIX_TOOL_FUNCTIONS** "`pmix.tool.fns`" (`bool`)

16 Request a list of functions supported by the PMIx tool library.

17 **PMIX_TOOL_ATTRIBUTES** "`pmix.setup.env`" (`bool`)

18 Request attributes supported by the PMIx tool library functions.

19 Server attributes

20 **PMIX_TOPOLOGY2** "`pmix.topo2`" (`pmix_topology_t`)

21 Provide a pointer to an implementation-specific description of the local node topology.

22 **PMIX_SERVER_SHARE_TOPOLOGY** "`pmix.srvr.share`" (`bool`)

23 The PMIx server is to share its copy of the local node topology (whether given to it or
24 self-discovered) with any clients.

25 **PMIX_SERVER_SESSION_SUPPORT** "`pmix.srvr.sess`" (`bool`)

26 The host RM wants to declare itself as being the local session server for PMIx connection
27 requests.

28 **PMIX_SERVER_START_TIME** "`pmix.srvr.strtime`" (`char*`)

29 Time when the server started - i.e., when the server created its rendezvous file (given in
30 ctime string format).

31 **PMIX_SERVER_SCHEDULER** "`pmix.srv.sched`" (`bool`)

32 Server is supporting system scheduler and desires access to appropriate WLM-supporting
33 features. Indicates that the library is to be initialized for scheduler support.

34 **PMIX_JOB_INFO_ARRAY** "`pmix.job.arr`" (`pmix_data_array_t`)

1 Provide an array of `pmix_info_t` containing job-realm information. The
2 `PMIX_SESSION_ID` attribute of the *session* containing the *job* is required to be included in
3 the array whenever the PMIx server library may host multiple sessions (e.g., when executing
4 with a host RM daemon). As information is registered one job (aka namespace) at a time via
5 the `PMIx_server_register_namespace` API, there is no requirement that the array
6 contain either the `PMIX_NAMESPACE` or `PMIX_JOBID` attributes when used in that context
7 (though either or both of them may be included). At least one of the job identifiers must be
8 provided in all other contexts where the job being referenced is ambiguous.

9 **`PMIX_APP_INFO_ARRAY`** "`pmix.app.arr`" (`pmix_data_array_t`)

10 Provide an array of `pmix_info_t` containing application-realm information. The
11 `PMIX_NAMESPACE` or `PMIX_JOBID` attributes of the *job* containing the application, plus its
12 `PMIX_APPNUM` attribute, must to be included in the array when the array is *not* included as
13 part of a call to `PMIx_server_register_namespace` - i.e., when the job containing the
14 application is ambiguous. The job identification is otherwise optional.

15 **`PMIX_PROC_INFO_ARRAY`** "`pmix.pdata`" (`pmix_data_array_t`)

16 Provide an array of `pmix_info_t` containing process-realm information. The
17 `PMIX_RANK` and `PMIX_NAMESPACE` attributes, or the `PMIX_PROCID` attribute, are required
18 to be included in the array when the array is not included as part of a call to
19 `PMIx_server_register_namespace` - i.e., when the job containing the process is
20 ambiguous. All three may be included if desired. When the array is included in some
21 broader structure that identifies the job, then only the `PMIX_RANK` or the `PMIX_PROCID`
22 attribute must be included (the others are optional).

23 **`PMIX_NODE_INFO_ARRAY`** "`pmix.node.arr`" (`pmix_data_array_t`)

24 Provide an array of `pmix_info_t` containing node-realm information. At a minimum,
25 either the `PMIX_NODEID` or `PMIX_HOSTNAME` attribute is required to be included in the
26 array, though both may be included.

27 **`PMIX_MAX_VALUE`** "`pmix.descr.maxval`" (*varies*)

28 Used in `pmix_regattr_t` to describe the maximum valid value for the associated
29 attribute.

30 **`PMIX_MIN_VALUE`** "`pmix.descr.minval`" (*varies*)

31 Used in `pmix_regattr_t` to describe the minimum valid value for the associated
32 attribute.

33 **`PMIX_ENUM_VALUE`** "`pmix.descr.enum`" (`char*`)

34 Used in `pmix_regattr_t` to describe accepted values for the associated attribute.
35 Numerical values shall be presented in a form convertible to the attribute's declared data
36 type. Named values (i.e., values defined by constant names via a typical C-language enum
37 declaration) must be provided as their numerical equivalent.

38 **`PMIX_HOMOGENEOUS_SYSTEM`** "`pmix.homo`" (`bool`)

39 The nodes comprising the session are homogeneous - i.e., they each contain the same
40 number of identical packages, fabric interfaces, GPUs, and other devices.

1 **PMIX_REQUIRED_KEY** "pmix.req.key" (char*)

2 Identifies a key that must be included in the requested information. If the specified key is not
3 already available, then the PMIx servers are required to delay response to the dmodex
4 request until either the key becomes available or the request times out.

5 Job-Mgmt attributes

6 **PMIX_ALLOC_ID** "pmix.alloc.id" (char*)

7 A string identifier (provided by the host environment) for the resulting allocation which can
8 later be used to reference the allocated resources in, for example, a call to **PMIx_Spawn**.

9 **PMIX_ALLOC_QUEUE** "pmix.alloc.queue" (char*)

10 Name of the WLM queue to which the allocation request is to be directed, or the queue being
11 referenced in a query.

12 Publish attributes

13 **PMIX_ACCESS_PERMISSIONS** "pmix.apperms" (pmix_data_array_t)

14 Define access permissions for the published data. The value shall contain an array of
15 **pmix_info_t** structs containing the specified permissions.

16 **PMIX_ACCESS_USERIDS** "pmix.auids" (pmix_data_array_t)

17 Array of effective UIDs that are allowed to access the published data.

18 **PMIX_ACCESS_GRPIDS** "pmix.agids" (pmix_data_array_t)

19 Array of effective GIDs that are allowed to access the published data.

20 Reserved keys

21 **PMIX_NUM_ALLOCATED_NODES** "pmix.num.anodes" (uint32_t)

22 Number of nodes in the specified realm regardless of whether or not they currently host
23 processes. Defaults to the *job* realm.

24 **PMIX_NUM_NODES** "pmix.num.nodes" (uint32_t)

25 Number of nodes currently hosting processes in the specified realm. Defaults to the *job*
26 realm.

27 **PMIX_CMD_LINE** "pmix.cmd.line" (char*)

28 Command line used to execute the specified job (e.g., "mpirun -n 2 -map-by foo ./myapp : -n
29 4 ./myapp2"). If the job was created by a call to **PMIx_Spawn**, the string is an inorder
30 concatenation of the values of **PMIX_APP_ARGV** for each application in the job using the
31 character ':' as a separator.

32 **PMIX_APP_ARGV** "pmix.app.argv" (char*)

33 Consolidated argv passed to the spawn command for the given application (e.g., "./myapp
34 arg1 arg2 arg3").

35 **PMIX_PACKAGE_RANK** "pmix.pkgrank" (uint16_t)

1 Rank of the specified process on the *package* where this process resides - refers to the
2 numerical location (starting from zero) of the process on its package when counting only
3 those processes from the same job that share the package, ordered by their overall rank
4 within that job. Note that processes that are not bound to PUs within a single specific
5 package cannot have a package rank.

6 **PMIX_REINCARNATION** "pmix.reinc" (uint32_t)

7 Number of times this process has been re-instantiated - i.e, a value of zero indicates that the
8 process has never been restarted.

9 **PMIX_HOSTNAME_ALIASES** "pmix.alias" (char*)

10 Comma-delimited list of names by which the target node is known.

11 **PMIX_HOSTNAME_KEEP_FQDN** "pmix.fqdn" (bool)

12 FQDNs are being retained by the PMIx library.

13 **PMIX_CPUSSET_BITMAP** "pmix.bitmap" (pmix_cpuset_t*)

14 Bitmap applied to the process upon launch.

15 **PMIX_EXTERNAL_PROGRESS** "pmix.evext" (bool)

16 The host shall progress the PMIx library via calls to **PMIx_Progress**

17 **PMIX_NODE_MAP_RAW** "pmix.nmap.raw" (char*)

18 Comma-delimited list of nodes containing procs within the specified realm. Defaults to the
19 *job* realm.

20 **PMIX_PROC_MAP_RAW** "pmix.pmap.raw" (char*)

21 Semi-colon delimited list of strings, each string containing a comma-delimited list of ranks
22 on the corresponding node within the specified realm. Defaults to the *job* realm.

23 Tool attributes

24 **PMIX_TOOL_CONNECT_OPTIONAL** "pmix.tool.conopt" (bool)

25 The tool shall connect to a server if available, but otherwise continue to operate
26 unconnected.

27 **PMIX_TOOL_ATTACHMENT_FILE** "pmix.tool.attach" (char*)

28 Pathname of file containing connection information to be used for attaching to a specific
29 server.

30 **PMIX_LAUNCHER_RENDEZVOUS_FILE** "pmix.tool.lncrnd" (char*)

31 Pathname of file where the launcher is to store its connection information so that the
32 spawning tool can connect to it.

33 **PMIX_PRIMARY_SERVER** "pmix.pri.srvr" (bool)

34 The server to which the tool is connecting shall be designated the *primary* server once
35 connection has been accomplished.

36 **PMIX_NOHUP** "pmix.nohup" (bool)

1 Any processes started on behalf of the calling tool (or the specified namespace, if such
2 specification is included in the list of attributes) should continue after the tool disconnects
3 from its server.

4 **PMIX_LAUNCHER_DAEMON** "pmix.lnch.dmn" (char*)

5 Path to executable that is to be used as the backend daemon for the launcher. This replaces
6 the launcher's own daemon with the specified executable. Note that the user is therefore
7 responsible for ensuring compatibility of the specified executable and the host launcher.

8 **PMIX_FORKEXEC_AGENT** "pmix.frkex.agnt" (char*)

9 Path to executable that the launcher's backend daemons are to fork/exec in place of the actual
10 application processes. The fork/exec agent shall connect back (as a PMIx tool) to the
11 launcher's daemon to receive its spawn instructions, and is responsible for starting the actual
12 application process it replaced. See Section 18.4.3 for details.

13 **PMIX_EXEC_AGENT** "pmix.exec.agnt" (char*)

14 Path to executable that the launcher's backend daemons are to fork/exec in place of the actual
15 application processes. The launcher's daemon shall pass the full command line of the
16 application on the command line of the exec agent, which shall not connect back to the
17 launcher's daemon. The exec agent is responsible for exec'ing the specified application
18 process in its own place. See Section 18.4.3 for details.

19 **PMIX_IOF_PUSH_STDIN** "pmix.iof.stdin" (bool)

20 Requests that the PMIx library collect the `stdin` of the requester and forward it to the
21 processes specified in the `PMIx_IOF_push` call. All collected data is sent to the same
22 targets until `stdin` is closed, or a subsequent call to `PMIx_IOF_push` is made that
23 includes the `PMIX_IOF_COMPLETE` attribute indicating that forwarding of `stdin` is to be
24 terminated.

25 **PMIX_IOF_COPY** "pmix.iof.cpy" (bool)

26 Requests that the host environment deliver a copy of the specified output stream(s) to the
27 tool, letting the stream(s) continue to also be delivered to the default location. This allows the
28 tool to tap into the output stream(s) without redirecting it from its current final destination.

29 **PMIX_IOF_REDIRECT** "pmix.iof.redir" (bool)

30 Requests that the host environment intercept the specified output stream(s) and deliver it to
31 the requesting tool instead of its current final destination. This might be used, for example,
32 during a debugging procedure to avoid injection of debugger-related output into the
33 application's results file. The original output stream(s) destination is restored upon
34 termination of the tool.

35 **PMIX_DEBUG_TARGET** "pmix.dbg.tgt" (pmix_proc_t*)

36 Identifier of process(es) to be debugged - a rank of `PMIX_RANK_WILDCARD` indicates that
37 all processes in the specified namespace are to be included.

38 **PMIX_DEBUG_DAEMONS_PER_PROC** "pmix.dbg.dpproc" (uint16_t)

1 Number of debugger daemons to be spawned per application process. The launcher is to pass
2 the identifier of the namespace to be debugged by including the **PMIX_DEBUG_TARGET**
3 attribute in the daemon's job-level information. The debugger daemons spawned on a given
4 node are responsible for self-determining their specific target process(es) - e.g., by
5 referencing their own **PMIX_LOCAL_RANK** in the daemon debugger job versus the
6 corresponding **PMIX_LOCAL_RANK** of the target processes on the node.

7 **PMIX_DEBUG_DAEMONS_PER_NODE** "pmix.dbg.dpnd" (uint16_t)

8 Number of debugger daemons to be spawned on each node where the target job is executing.
9 The launcher is to pass the identifier of the namespace to be debugged by including the
10 **PMIX_DEBUG_TARGET** attribute in the daemon's job-level information. The debugger
11 daemons spawned on a given node are responsible for self-determining their specific target
12 process(es) - e.g., by referencing their own **PMIX_LOCAL_RANK** in the daemon debugger
13 job versus the corresponding **PMIX_LOCAL_RANK** of the target processes on the node.

14 **PMIX_WAIT_FOR_CONNECTION** "pmix.wait.conn" (bool)

15 Wait until the specified process has connected to the requesting tool or server, or the
16 operation times out (if the **PMIX_TIMEOUT** directive is included in the request).

17 **PMIX_LAUNCH_DIRECTIVES** "pmix.lnch.dirs" (pmix_data_array_t*)

18 Array of **pmix_info_t** containing directives for the launcher - a convenience attribute for
19 retrieving all directives with a single call to **PMIx_Get**.

20 Fabric attributes

21 **PMIX_SERVER_SCHEDULER** "pmix.srv.sched" (bool)

22 Server is supporting system scheduler and desires access to appropriate WLM-supporting
23 features. Indicates that the library is to be initialized for scheduler support.

24 **PMIX_FABRIC_COST_MATRIX** "pmix.fab.cm" (pointer)

25 Pointer to a two-dimensional square array of point-to-point relative communication costs
26 expressed as **uint16_t** values.

27 **PMIX_FABRIC_GROUPS** "pmix.fab.grps" (string)

28 A string delineating the group membership of nodes in the overall system, where each fabric
29 group consists of the group number followed by a colon and a comma-delimited list of nodes
30 in that group, with the groups delimited by semi-colons (e.g.,

31 0:node000,node002,node004,node006;1:node001,node003,
32 node005,node007)

33 **PMIX_FABRIC_VENDOR** "pmix.fab.vndr" (string)

34 Name of the vendor (e.g., Amazon, Mellanox, HPE, Intel) for the specified fabric.

35 **PMIX_FABRIC_IDENTIFIER** "pmix.fab.id" (string)

36 An identifier for the specified fabric (e.g., MgmtEthernet, Slingshot-11, OmniPath-1).

37 **PMIX_FABRIC_INDEX** "pmix.fab.idx" (size_t)

38 The index of the fabric as returned in **pmix_fabric_t**.

1 **PMIX_FABRIC_NUM_DEVICES** "pmix.fab.nverts" (**size_t**)
2 Total number of fabric devices in the overall system - corresponds to the number of rows or
3 columns in the cost matrix.

4 **PMIX_FABRIC_COORDINATES** "pmix.fab.coords" (**pmix_data_array_t**)
5 Array of **pmix_geometry_t** fabric coordinates for devices on the specified node. The
6 array will contain the coordinates of all devices on the node, including values for all
7 supported coordinate views. The information for devices on the local node shall be provided
8 if the node is not specified in the request.

9 **PMIX_FABRIC_DIMS** "pmix.fab.dims" (**uint32_t**)
10 Number of dimensions in the specified fabric plane/view. If no plane is specified in a
11 request, then the dimensions of all planes in the overall system will be returned as a
12 **pmix_data_array_t** containing an array of **uint32_t** values. Default is to provide
13 dimensions in *logical* view.

14 **PMIX_FABRIC_ENDPT** "pmix.fab.endpt" (**pmix_data_array_t**)
15 Fabric endpoints for a specified process. As multiple endpoints may be assigned to a given
16 process (e.g., in the case where multiple devices are associated with a package to which the
17 process is bound), the returned values will be provided in a **pmix_data_array_t** of
18 **pmix_endpoint_t** elements.

19 **PMIX_FABRIC_SHAPE** "pmix.fab.shape" (**pmix_data_array_t***)
20 The size of each dimension in the specified fabric plane/view, returned in a
21 **pmix_data_array_t** containing an array of **uint32_t** values. The size is defined as
22 the number of elements present in that dimension - e.g., the number of devices in one
23 dimension of a physical view of a fabric plane. If no plane is specified, then the shape of
24 each plane in the overall system will be returned in a **pmix_data_array_t** array where
25 each element is itself a two-element array containing the **PMIX_FABRIC_PLANE** followed
26 by that plane's fabric shape. Default is to provide the shape in *logical* view.

27 **PMIX_FABRIC_SHAPE_STRING** "pmix.fab.shapestr" (**string**)
28 Network shape expressed as a string (e.g., "10x12x2"). If no plane is specified, then the
29 shape of each plane in the overall system will be returned in a **pmix_data_array_t** array
30 where each element is itself a two-element array containing the **PMIX_FABRIC_PLANE**
31 followed by that plane's fabric shape string. Default is to provide the shape in *logical* view.

32 **PMIX_SWITCH_PEERS** "pmix.speers" (**pmix_data_array_t**)
33 Peer ranks that share the same switch as the process specified in the call to **PMIx_Get**.
34 Returns a **pmix_data_array_t** array of **pmix_info_t** results, each element
35 containing the **PMIX_SWITCH_PEERS** key with a three-element **pmix_data_array_t**
36 array of **pmix_info_t** containing the **PMIX_DEVICE_ID** of the local fabric device, the
37 **PMIX_FABRIC_SWITCH** identifying the switch to which it is connected, and a
38 comma-delimited string of peer ranks sharing the switch to which that device is connected.

39 **PMIX_FABRIC_PLANE** "pmix.fab.plane" (**string**)

1 ID string of a fabric plane (e.g., CIDR for Ethernet). When used as a modifier in a request
2 for information, specifies the plane whose information is to be returned. When used directly
3 as a key in a request, returns a `pmix_data_array_t` of string identifiers for all fabric
4 planes in the overall system.

5 **`PMIX_FABRIC_SWITCH`** "`pmix.fab.switch`" (`string`)

6 ID string of a fabric switch. When used as a modifier in a request for information, specifies
7 the switch whose information is to be returned. When used directly as a key in a request,
8 returns a `pmix_data_array_t` of string identifiers for all fabric switches in the overall
9 system.

10 **`PMIX_FABRIC_DEVICE`** "`pmix.fabdev`" (`pmix_data_array_t`)

11 An array of `pmix_info_t` describing a particular fabric device using one or more of the
12 attributes defined below. The first element in the array shall be the `PMIX_DEVICE_ID` of
13 the device.

14 **`PMIX_FABRIC_DEVICE_INDEX`** "`pmix.fabdev.idx`" (`uint32_t`)

15 Index of the device within an associated communication cost matrix.

16 **`PMIX_FABRIC_DEVICE_NAME`** "`pmix.fabdev.nm`" (`string`)

17 The operating system name associated with the device. This may be a logical fabric interface
18 name (e.g. "eth0" or "eno1") or an absolute filename.

19 **`PMIX_FABRIC_DEVICE_VENDOR`** "`pmix.fabdev.vndr`" (`string`)

20 Indicates the name of the vendor that distributes the device.

21 **`PMIX_FABRIC_DEVICE_BUS_TYPE`** "`pmix.fabdev.btyp`" (`string`)

22 The type of bus to which the device is attached (e.g., "PCI", "GEN-Z").

23 **`PMIX_FABRIC_DEVICE_VENDORID`** "`pmix.fabdev.vendid`" (`string`)

24 This is a vendor-provided identifier for the device or product.

25 **`PMIX_FABRIC_DEVICE_DRIVER`** "`pmix.fabdev.driver`" (`string`)

26 The name of the driver associated with the device.

27 **`PMIX_FABRIC_DEVICE_FIRMWARE`** "`pmix.fabdev.fmwr`" (`string`)

28 The device's firmware version.

29 **`PMIX_FABRIC_DEVICE_ADDRESS`** "`pmix.fabdev.addr`" (`string`)

30 The primary link-level address associated with the device, such as a MAC address. If
31 multiple addresses are available, only one will be reported.

32 **`PMIX_FABRIC_DEVICE_COORDINATES`** "`pmix.fab.coord`" (`pmix_geometry_t`)

33 The `pmix_geometry_t` fabric coordinates for the device, including values for all
34 supported coordinate views.

35 **`PMIX_FABRIC_DEVICE_MTU`** "`pmix.fabdev.mtu`" (`size_t`)

36 The maximum transfer unit of link level frames or packets, in bytes.

37 **`PMIX_FABRIC_DEVICE_SPEED`** "`pmix.fabdev.speed`" (`size_t`)

1 The active link data rate, given in bits per second.

2 **PMIX_FABRIC_DEVICE_STATE** "pmix.fabdev.state" (**pmix_link_state_t**)

3 The last available physical port state for the specified device. Possible values are
4 **PMIX_LINK_STATE_UNKNOWN**, **PMIX_LINK_DOWN**, and **PMIX_LINK_UP**, to indicate
5 if the port state is unknown or not applicable (unknown), inactive (down), or active (up).

6 **PMIX_FABRIC_DEVICE_TYPE** "pmix.fabdev.type" (**string**)

7 Specifies the type of fabric interface currently active on the device, such as Ethernet or
8 InfiniBand.

9 **PMIX_FABRIC_DEVICE_PCI_DEVID** "pmix.fabdev.pcidevid" (**string**)

10 A node-level unique identifier for a PCI device. Provided only if the device is located on a
11 PCI bus. The identifier is constructed as a four-part tuple delimited by colons comprised of
12 the PCI 16-bit domain, 8-bit bus, 8-bit device, and 8-bit function IDs, each expressed in
13 zero-extended hexadecimal form. Thus, an example identifier might be "abc1:0f:23:01". The
14 combination of node identifier (**PMIX_HOSTNAME** or **PMIX_NODEID**) and
15 **PMIX_FABRIC_DEVICE_PCI_DEVID** shall be unique within the overall system.

16 Device attributes

17 **PMIX_DEVICE_DISTANCES** "pmix.dev.dist" (**pmix_data_array_t**)

18 Return an array of **pmix_device_distance_t** containing the minimum and maximum
19 distances of the given process location to all devices of the specified type on the local node.

20 **PMIX_DEVICE_TYPE** "pmix.dev.type" (**pmix_device_type_t**)

21 Bitmask specifying the type(s) of device(s) whose information is being requested. Only used
22 as a directive/qualifier.

23 **PMIX_DEVICE_ID** "pmix.dev.id" (**string**)

24 System-wide UUID or node-local OS name of a particular device.

25 Sets-Groups attributes

26 **PMIX_QUERY_NUM_PSETS** "pmix.qry.psetnum" (**size_t**)

27 Return the number of process sets defined in the specified range (defaults to
28 **PMIX_RANGE_SESSION**).

29 **PMIX_QUERY_PSET_NAMES** "pmix.qry.psets" (**pmix_data_array_t***)

30 Return a **pmix_data_array_t** containing an array of strings of the process set names
31 defined in the specified range (defaults to **PMIX_RANGE_SESSION**).

32 **PMIX_QUERY_PSET_MEMBERSHIP** "pmix.qry.pmems" (**pmix_data_array_t***)

33 Return an array of **pmix_proc_t** containing the members of the specified process set.

34 **PMIX_PSET_NAME** "pmix.pset.nm" (**char***)

35 The name of the newly defined process set.

36 **PMIX_PSET_MEMBERS** "pmix.pset.mems" (**pmix_data_array_t***)

37 An array of **pmix_proc_t** containing the members of the newly defined process set.

1 **PMIX_PSET_NAMES** "pmix.pset.nms" (pmix_data_array_t*)
2 Returns an array of **char*** string names of the process sets in which the given process is a
3 member.

4 **PMIX_QUERY_NUM_GROUPS** "pmix.qry.pgrpnum" (size_t)
5 Return the number of process groups defined in the specified range (defaults to session).
6 OPTIONAL QUALIFERS: **PMIX_RANGE**.

7 **PMIX_QUERY_GROUP_NAMES** "pmix.qry.pgrp" (pmix_data_array_t*)
8 Return a **pmix_data_array_t** containing an array of string names of the process groups
9 defined in the specified range (defaults to session). OPTIONAL QUALIFERS:
10 **PMIX_RANGE**.

11 **PMIX_QUERY_GROUP_MEMBERSHIP**
12 "pmix.qry.pgrpmems" (pmix_data_array_t*)
13 Return a **pmix_data_array_t** of **pmix_proc_t** containing the members of the
14 specified process group. REQUIRED QUALIFIERS: **PMIX_GROUP_ID**.

15 **PMIX_GROUP_ID** "pmix.grp.id" (char*)
16 User-provided group identifier - as the group identifier may be used in PMIx operations, the
17 user is required to ensure that the provided ID is unique within the scope of the host
18 environment (e.g., by including some user-specific or application-specific prefix or suffix to
19 the string).

20 **PMIX_GROUP_LEADER** "pmix.grp.ldr" (bool)
21 This process is the leader of the group.

22 **PMIX_GROUP_OPTIONAL** "pmix.grp.opt" (bool)
23 Participation is optional - do not return an error if any of the specified processes terminate
24 without having joined. The default is **false**.

25 **PMIX_GROUP_NOTIFY_TERMINATION** "pmix.grp.notterm" (bool)
26 Notify remaining members when another member terminates without first leaving the group.
27

28 **PMIX_GROUP_FT_COLLECTIVE** "pmix.grp.ftcoll" (bool)
29 Adjust internal tracking on-the-fly for terminated processes during a PMIx group collective
30 operation.

31 **PMIX_GROUP_ASSIGN_CONTEXT_ID** "pmix.grp.actxid" (bool)
32 Requests that the RM assign a new context identifier to the newly created group. The
33 identifier is an unsigned, **size_t** value that the RM guarantees to be unique across the range
34 specified in the request. Thus, the value serves as a means of identifying the group within
35 that range. If no range is specified, then the request defaults to **PMIX_RANGE_SESSION**.

36 **PMIX_GROUP_LOCAL_ONLY** "pmix.grp.lcl" (bool)
37 Group operation only involves local processes. PMIx implementations are *required* to
38 automatically scan an array of group members for local vs remote processes - if only local

1 processes are detected, the implementation need not execute a global collective for the
2 operation unless a context ID has been requested from the host environment. This can result
3 in significant time savings. This attribute can be used to optimize the operation by indicating
4 whether or not only local processes are represented, thus allowing the implementation to
5 bypass the scan.

6 **PMIX_GROUP_CONTEXT_ID** "pmix.grp.ctxid" (size_t)

7 Context identifier assigned to the group by the host RM.

8 **PMIX_GROUP_ENDPT_DATA** "pmix.grp.endpt" (pmix_byte_object_t)

9 Data collected during group construction to ensure communication between group members
10 is supported upon completion of the operation.

11 **PMIX_GROUP_NAMES** "pmix.pgrp.nm" (pmix_data_array_t*)

12 Returns an array of **char*** string names of the process groups in which the given process is
13 a member.

14 Process Mgmt attributes

15 **PMIX_OUTPUT_TO_DIRECTORY** "pmix.outdir" (char*)

16 Direct output into files of form "<directory>/<jobid>/rank.<rank>/
17 **stdout[err]**" - can be assigned to the entire job (by including attribute in the *job_info*
18 array) or on a per-application basis in the *info* array for each **pmix_app_t**.

19 **PMIX_TIMEOUT_STACKTRACES** "pmix.tim.stack" (bool)

20 Include process stacktraces in timeout report from a job.

21 **PMIX_TIMEOUT_REPORT_STATE** "pmix.tim.state" (bool)

22 Report process states in timeout report from a job.

23 **PMIX_NOTIFY_JOB_EVENTS** "pmix.note.jev" (bool)

24 Requests that the launcher generate the **PMIX_EVENT_JOB_START**,
25 **PMIX_LAUNCH_COMPLETE**, and **PMIX_EVENT_JOB_END** events. Each event is to
26 include at least the namespace of the corresponding job and a **PMIX_EVENT_TIMESTAMP**
27 indicating the time the event occurred. Note that the requester must register for these
28 individual events, or capture and process them by registering a default event handler instead
29 of individual handlers and then process the events based on the returned status code. Another
30 common method is to register one event handler for all job-related events, with a separate
31 handler for non-job events - see **PMIx_Register_event_handler** for details.

32 **PMIX_NOTIFY_PROC_TERMINATION** "pmix.noteproc" (bool)

33 Requests that the launcher generate the **PMIX_EVENT_PROC_TERMINATED** event
34 whenever a process either normally or abnormally terminates.

35 **PMIX_NOTIFY_PROC_ABNORMAL_TERMINATION** "pmix.noteabproc" (bool)

36 Requests that the launcher generate the **PMIX_EVENT_PROC_TERMINATED** event only
37 when a process abnormally terminates.

38 **PMIX_LOG_PROC_TERMINATION** "pmix.logproc" (bool)

1 Requests that the launcher log the `PMIX_EVENT_PROC_TERMINATED` event whenever a
2 process either normally or abnormally terminates.

3 `PMIX_LOG_PROC_ABNORMAL_TERMINATION` "pmix.logabproc" (bool)

4 Requests that the launcher log the `PMIX_EVENT_PROC_TERMINATED` event only when a
5 process abnormally terminates.

6 `PMIX_LOG_JOB_EVENTS` "pmix.log.jev" (bool)

7 Requests that the launcher log the `PMIX_EVENT_JOB_START`,
8 `PMIX_LAUNCH_COMPLETE`, and `PMIX_EVENT_JOB_END` events using `PMIx_Log`,
9 subject to the logging attributes of Section 13.4.3.

10 `PMIX_LOG_COMPLETION` "pmix.logcomp" (bool)

11 Requests that the launcher log the `PMIX_EVENT_JOB_END` event for normal or abnormal
12 termination of the spawned job using `PMIx_Log`, subject to the logging attributes of
13 Section 13.4.3. The event shall include the returned status code
14 (`PMIX_JOB_TERM_STATUS`) for the corresponding job; the identity (`PMIX_PROCID`)
15 and exit status (`PMIX_EXIT_CODE`) of the first failed process, if applicable; and a
16 `PMIX_EVENT_TIMESTAMP` indicating the time the termination occurred.

17 `PMIX_FIRST_ENVAR` "pmix.envar.first" (pmix_envar_t*)

18 Ensure the given value appears first in the specified envar using the separator character,
19 creating the envar if it doesn't already exist

20 Event attributes

21 `PMIX_EVENT_TIMESTAMP` "pmix.evtstamp" (time_t)

22 System time when the associated event occurred.

23 C.8.3 Added Environmental Variables

24 Tool environmental variables

25 `PMIX_LAUNCHER_RNDZ_URI`

26 `PMIX_LAUNCHER_RNDZ_FILE`

27 `PMIX_KEEPA_LIVE_PIPE`

29 C.8.4 Added Macros

30 `PMIX_CHECK_RESERVED_KEY` `PMIX_INFO_WAS_PROCESSED` `PMIX_INFO_PROCESSED`

31 `PMIX_INFO_LIST_START` `PMIX_INFO_LIST_ADD` `PMIX_INFO_LIST_XFER`

32 `PMIX_INFO_LIST_CONVERT` `PMIX_INFO_LIST_RELEASE`

1 C.8.5 Deprecated APIs

2 `pmix_evhdlr_reg_cbfunc_t` Renamed to `pmix_hdlr_reg_cbfunc_t`

3 The `pmix_server_client_connected_fn_t` server module entry point has been
4 *deprecated* in favor of `pmix_server_client_connected2_fn_t`

5 `PMix_tool_connect_to_server` Replaced by `PMix_tool_attach_to_server` to
6 allow return of the process identifier of the server to which the tool has attached.

7 C.8.6 Deprecated constants

8 The following constants were deprecated in v4.0:

9 `PMIX_ERR_DEBUGGER_RELEASE` Renamed to `PMIX_DEBUGGER_RELEASE`
10 `PMIX_ERR_JOB_TERMINATED` Renamed to `PMIX_EVENT_JOB_END`
11 `PMIX_EXISTS` Renamed to `PMIX_ERR_EXISTS`
12 `PMIX_ERR_PROC_ABORTED` Consolidated with `PMIX_EVENT_PROC_TERMINATED`
13 `PMIX_ERR_PROC_ABORTING` Consolidated with `PMIX_EVENT_PROC_TERMINATED`
14 `PMIX_ERR_LOST_CONNECTION_TO_SERVER` Consolidated into
15 `PMIX_ERR_LOST_CONNECTION`
16 `PMIX_ERR_LOST_PEER_CONNECTION` Consolidated into
17 `PMIX_ERR_LOST_CONNECTION`
18 `PMIX_ERR_LOST_CONNECTION_TO_CLIENT` Consolidated into
19 `PMIX_ERR_LOST_CONNECTION`
20 `PMIX_ERR_INVALID_TERMINATION` Renamed to `PMIX_ERR_JOB_TERM_WO_SYNC`
21 `PMIX_PROC_TERMINATED` Renamed to `PMIX_EVENT_PROC_TERMINATED`
22 `PMIX_ERR_NODE_DOWN` Renamed to `PMIX_EVENT_NODE_DOWN`
23 `PMIX_ERR_NODE_OFFLINE` Renamed to `PMIX_EVENT_NODE_OFFLINE`
24 `PMIX_ERR_SYS_OTHER` Renamed to `PMIX_EVENT_SYS_OTHER`
25 `PMIX_CONNECT_REQUESTED` Connection has been requested by a PMIx-based tool -
26 deprecated as not required.
27 `PMIX_PROC_HAS_CONNECTED` A tool or client has connected to the PMIx server -
28 deprecated in favor of the new `pmix_server_client_connected2_fn_t` server
29 module API

30 C.8.7 Removed constants

31 The following constants were removed from the PMIx Standard in v4.0 as they are internal to a
32 particular PMIx implementation.

33 `PMIX_ERR_HANDSHAKE_FAILED` Connection handshake failed
34 `PMIX_ERR_READY_FOR_HANDSHAKE` Ready for handshake
35 `PMIX_ERR_IN_ERRNO` Error defined in `errno`
36 `PMIX_ERR_INVALID_VAL_LENGTH` Invalid value length

1	PMIX_ERR_INVALID_LENGTH	Invalid argument length
2	PMIX_ERR_INVALID_NUM_ARGS	Invalid number of arguments
3	PMIX_ERR_INVALID_ARGS	Invalid arguments
4	PMIX_ERR_INVALID_NUM_PARSED	Invalid number parsed
5	PMIX_ERR_INVALID_KEYVALP	Invalid key/value pair
6	PMIX_ERR_INVALID_SIZE	Invalid size
7	PMIX_ERR_PROC_REQUESTED_ABORT	Process is already requested to abort
8	PMIX_ERR_SERVER_FAILED_REQUEST	Failed to connect to the server
9	PMIX_ERR_PROC_ENTRY_NOT_FOUND	Process not found
10	PMIX_ERR_INVALID_ARG	Invalid argument
11	PMIX_ERR_INVALID_KEY	Invalid key
12	PMIX_ERR_INVALID_KEY_LENGTH	Invalid key length
13	PMIX_ERR_INVALID_VAL	Invalid value
14	PMIX_ERR_INVALID_NAMESPACE	Invalid namespace
15	PMIX_ERR_SERVER_NOT_AVAIL	Server is not available
16	PMIX_ERR_SILENT	Silent error
17	PMIX_ERR_PACK_MISMATCH	Pack mismatch
18	PMIX_ERR_DATA_VALUE_NOT_FOUND	Data value not found
19	PMIX_ERR_NOT_IMPLEMENTED	Not implemented
20	PMIX_GDS_ACTION_COMPLETE	The GDS action has completed
21	PMIX_NOTIFY_ALLOC_COMPLETE	Notify that a requested allocation operation is complete
22		- the result of the request will be included in the <i>info</i> array

23 C.8.8 Deprecated attributes

24 The following attributes were deprecated in v4.0:

25	PMIX_TOPOLOGY	"pmix.topo" (hwloc_topology_t)
26		Renamed to PMIX_TOPOLOGY2 .
27	PMIX_DEBUG_JOB	"pmix.dbg.job" (char*)
28		Renamed to PMIX_DEBUG_TARGET
29	PMIX_RECONNECT_SERVER	"pmix.tool.recon" (bool)
30		Renamed to the PMIx_tool_connect_to_server API
31	PMIX_ALLOC_NETWORK	"pmix.alloc.net" (array)
32		Renamed to PMIX_ALLOC_FABRIC
33	PMIX_ALLOC_NETWORK_ID	"pmix.alloc.netid" (char*)
34		Renamed to PMIX_ALLOC_FABRIC_ID
35	PMIX_ALLOC_NETWORK_QOS	"pmix.alloc.netqos" (char*)
36		Renamed to PMIX_ALLOC_FABRIC_QOS
37	PMIX_ALLOC_NETWORK_TYPE	"pmix.alloc.nettype" (char*)
38		Renamed to PMIX_ALLOC_FABRIC_TYPE
39	PMIX_ALLOC_NETWORK_PLANE	"pmix.alloc.netplane" (char*)
40		Renamed to PMIX_ALLOC_FABRIC_PLANE

1 **PMIX_ALLOC_NETWORK_ENDPTS** "pmix.alloc.endpts" (size_t)
 2 Renamed to **PMIX_ALLOC_FABRIC_ENDPTS**
 3 **PMIX_ALLOC_NETWORK_ENDPTS_NODE** "pmix.alloc.endpts.nd" (size_t)
 4 Renamed to **PMIX_ALLOC_FABRIC_ENDPTS_NODE**
 5 **PMIX_ALLOC_NETWORK_SEC_KEY** "pmix.alloc.nsec" (pmix_byte_object_t)
 6 Renamed to **PMIX_ALLOC_FABRIC_SEC_KEY**
 7 **PMIX_PROC_DATA** "pmix.pdata" (pmix_data_array_t)
 8 Renamed to **PMIX_PROC_INFO_ARRAY**
 9 **PMIX_LOCALITY** "pmix.loc" (pmix_locality_t)
 10 Relative locality of the specified process to the requester, expressed as a bitmask as per the
 11 description in the **pmix_locality_t** section. This value is unique to the requesting
 12 process and thus cannot be communicated by the server as part of the job-level information.
 13 Its use has been replaced by the **PMIx_Get_relative_locality** function.

14 C.8.9 Removed attributes

15 The following attributes were removed from the PMIx Standard in v4.0 as they are internal to a
 16 particular PMIx implementation. Users are referred to the **PMIx_Load_topology** API for
 17 obtaining the local topology description.

18 **PMIX_LOCAL_TOPO** "pmix.ltopo" (char*)
 19 XML representation of local node topology.
 20 **PMIX_TOPOLOGY_XML** "pmix.topo.xml" (char*)
 21 XML-based description of topology
 22 **PMIX_TOPOLOGY_FILE** "pmix.topo.file" (char*)
 23 Full path to file containing XML topology description
 24 **PMIX_TOPOLOGY_SIGNATURE** "pmix.toposig" (char*)
 25 Topology signature string.
 26 **PMIX_HWLOC_SHMEM_ADDR** "pmix.hwlocaddr" (size_t)
 27 Address of the HWLOC shared memory segment.
 28 **PMIX_HWLOC_SHMEM_SIZE** "pmix.hwlocsize" (size_t)
 29 Size of the HWLOC shared memory segment.
 30 **PMIX_HWLOC_SHMEM_FILE** "pmix.hwlocfile" (char*)
 31 Path to the HWLOC shared memory file.
 32 **PMIX_HWLOC_XML_V1** "pmix.hwlocxml1" (char*)
 33 XML representation of local topology using HWLOC's v1.x format.
 34 **PMIX_HWLOC_XML_V2** "pmix.hwlocxml2" (char*)
 35 XML representation of local topology using HWLOC's v2.x format.
 36 **PMIX_HWLOC_SHARE_TOPO** "pmix.hwlocsh" (bool)
 37 Share the HWLOC topology via shared memory
 38 **PMIX_HWLOC_HOLE_KIND** "pmix.hwlocholek" (char*)
 39 Kind of VM "hole" HWLOC should use for shared memory
 40 **PMIX_DSTPATH** "pmix.dstpath" (char*)

1 Path to shared memory data storage (dstore) files. Deprecated from Standard as being
2 implementation specific.

3 **PMIX_COLLECTIVE_ALGO** "pmix.calgo" (char*)

4 Comma-delimited list of algorithms to use for the collective operation. PMIx does not
5 impose any requirements on a host environment's collective algorithms. Thus, the
6 acceptable values for this attribute will be environment-dependent - users are encouraged to
7 check their host environment for supported values.

8 **PMIX_COLLECTIVE_ALGO_REQD** "pmix.calreqd" (bool)

9 If **true**, indicates that the requested choice of algorithm is mandatory.

10 **PMIX_PROC_BLOB** "pmix.pblob" (pmix_byte_object_t)

11 Packed blob of process data.

12 **PMIX_MAP_BLOB** "pmix.mblob" (pmix_byte_object_t)

13 Packed blob of process location.

14 **PMIX_MAPPER** "pmix.mapper" (char*)

15 Mapping mechanism to use for placing spawned processes - when accessed using
16 [PMIx_Get](#), use the [PMIX_RANK_WILDCARD](#) value for the rank to discover the mapping
17 mechanism used for the provided namespace.

18 **PMIX_NON_PMI** "pmix.nonpmi" (bool)

19 Spawned processes will not call [PMIx_Init](#).

20 **PMIX_PROC_URI** "pmix.puri" (char*)

21 URI containing contact information for the specified process.

22 **PMIX_ARCH** "pmix.arch" (uint32_t)

23 Architecture flag.

24 C.9 Version 4.1: TBD

25 The v4.1 update includes clarifications and corrections from the v4.0 document:

26 C.9.1 Removed constants

27 The following constants were removed from the PMIx Standard in v4.1 as they are internal to a
28 particular PMIx implementation.

29 **PMIX_BUFFER** Buffer.

- 30 • Remove some stale language in [Chapter 9.1](#).
- 31 • Provisional Items:
 - 32 – Storage Chapter [19](#) on page [474](#)

33 C.9.2 Added Functions (Provisional)

- 34 • [PMIx_Data_load](#)
- 35 • [PMIx_Data_unload](#)
- 36 • [PMIx_Data_compress](#)
- 37 • [PMIx_Data_decompress](#)

1 C.9.3 Added Data Structures (Provisional)

- 2 • pmix_storage_medium_t
- 3 • pmix_storage_accessibility_t
- 4 • pmix_storage_persistence_t
- 5 • pmix_storage_access_type_t

6 C.9.4 Added Macros (Provisional)

- 7 • PMIX_NAMESPACE_INVALID
- 8 • PMIX_RANK_IS_VALID
- 9 • PMIX_PROCID_INVALID
- 10 • PMIX_PROCID_XFER

11 C.9.5 Added Constants (Provisional)

- 12 • PMIX_PROC_NAMESPACE

13 Storage constants

- 14 • PMIX_STORAGE_MEDIUM_UNKNOWN
- 15 • PMIX_STORAGE_MEDIUM_TAPE
- 16 • PMIX_STORAGE_MEDIUM_HDD
- 17 • PMIX_STORAGE_MEDIUM_SSD
- 18 • PMIX_STORAGE_MEDIUM_NVME
- 19 • PMIX_STORAGE_MEDIUM_PMEM
- 20 • PMIX_STORAGE_MEDIUM_RAM
- 21 • PMIX_STORAGE_ACCESSIBILITY_NODE
- 22 • PMIX_STORAGE_ACCESSIBILITY_SESSION
- 23 • PMIX_STORAGE_ACCESSIBILITY_JOB
- 24 • PMIX_STORAGE_ACCESSIBILITY_RACK
- 25 • PMIX_STORAGE_ACCESSIBILITY_CLUSTER
- 26 • PMIX_STORAGE_ACCESSIBILITY_REMOTE
- 27 • PMIX_STORAGE_PERSISTENCE_TEMPORARY
- 28 • PMIX_STORAGE_PERSISTENCE_NODE
- 29 • PMIX_STORAGE_PERSISTENCE_SESSION
- 30 • PMIX_STORAGE_PERSISTENCE_JOB
- 31 • PMIX_STORAGE_PERSISTENCE_SCRATCH
- 32 • PMIX_STORAGE_PERSISTENCE_PROJECT
- 33 • PMIX_STORAGE_PERSISTENCE_ARCHIVE
- 34 • PMIX_STORAGE_ACCESS_RD
- 35 • PMIX_STORAGE_ACCESS_WR
- 36 • PMIX_STORAGE_ACCESS_RDWR

1 C.9.6 Added Attributes (Provisional)

2 Storage attributes

3 **PMIX_STORAGE_ID** "pmix.strg.id" (char*)

4 An identifier for the storage system (e.g., lustre-fs1, daos-oss1, home-fs)

5 **PMIX_STORAGE_PATH** "pmix.strg.path" (char*)

6 Mount point path for the storage system (valid only for file-based storage systems)

7 **PMIX_STORAGE_TYPE** "pmix.strg.type" (char*)

8 Type of storage system (i.e., "lustre", "gpfs", "daos", "ext4")

9 **PMIX_STORAGE_VERSION** "pmix.strg.ver" (char*)

10 Version string for the storage system

11 **PMIX_STORAGE_MEDIUM** "pmix.strg.medium" (pmix_storage_medium_t)

12 Types of storage mediums utilized by the storage system (e.g., SSDs, HDDs, tape)

13 **PMIX_STORAGE_ACCESSIBILITY**

14 "pmix.strg.access" (pmix_storage_accessibility_t)

15 Accessibility level of the storage system (e.g., within same node, within same session)

16 **PMIX_STORAGE_PERSISTENCE**

17 "pmix.strg.persist" (pmix_storage_persistence_t)

18 Persistence level of the storage system (e.g., scratch storage or archive storage)

19 **PMIX_QUERY_STORAGE_LIST** "pmix.strg.list" (char*)

20 Comma-delimited list of storage identifiers (i.e., **PMIX_STORAGE_ID** types) for available
21 storage systems

22 **PMIX_STORAGE_CAPACITY_LIMIT** "pmix.strg.caplim" (double)

23 Overall limit on capacity (in bytes) for the storage system

24 **PMIX_STORAGE_CAPACITY_USED** "pmix.strg.capuse" (double)

25 Overall used capacity (in bytes) for the storage system

26 **PMIX_STORAGE_OBJECT_LIMIT** "pmix.strg.objlim" (uint64_t)

27 Overall limit on number of objects (e.g., inodes) for the storage system

28 **PMIX_STORAGE_OBJECTS_USED** "pmix.strg.objuse" (uint64_t)

29 Overall used number of objects (e.g., inodes) for the storage system

30 **PMIX_STORAGE_MINIMAL_XFER_SIZE** "pmix.strg.minxfer" (double)

31 Minimal transfer size (in bytes) for the storage system - this is the storage system's atomic
32 unit of transfer (e.g., block size)

33 **PMIX_STORAGE_SUGGESTED_XFER_SIZE** "pmix.strg.sxfer" (double)

34 Suggested transfer size (in bytes) for the storage system

35 **PMIX_STORAGE_BW_MAX** "pmix.strg.bwmax" (double)

1 Maximum bandwidth (in bytes/sec) for storage system - provided as the theoretical
2 maximum or the maximum observed bandwidth value

3 **PMIX_STORAGE_BW_CUR** "pmix.strg.bwcur" (double)

4 Observed bandwidth (in bytes/sec) for storage system - provided as a recently observed
5 bandwidth value, with the exact measurement interval depending on the storage system
6 and/or PMIx library implementation

7 **PMIX_STORAGE_IOPS_MAX** "pmix.strg.iopsmax" (double)

8 Maximum IOPS (in I/O operations per second) for storage system - provided as the
9 theoretical maximum or the maximum observed IOPS value

10 **PMIX_STORAGE_IOPS_CUR** "pmix.strg.iopscur" (double)

11 Observed IOPS (in I/O operations per second) for storage system - provided as a recently
12 observed IOPS value, with the exact measurement interval depending on the storage system
13 and/or PMIx library implementation

14 **PMIX_STORAGE_ACCESS_TYPE**

15 "pmix.strg.atype" (pmix_storage_access_type_t)

16 Qualifier describing the type of storage access to return information for (e.g., for qualifying
17 **PMIX_STORAGE_BW_CUR**, **PMIX_STORAGE_IOPS_CUR**, or
18 **PMIX_STORAGE_SUGGESTED_XFER_SIZE** attributes)

19 C.10 Version 4.2: TBD

20 The v4.2 update includes the following changes from the v4.1 document:

- 21 • Define when **PMIX_PARENT_ID** is set
- 22 • Add a definition for *tool*
- 23 • Clarify **PMIX_CMD_LINE** in **PMIx_Spawn**

24 C.10.1 Deprecated constants

25 The following constants were deprecated in v4.2:

26 **PMIX_DEBUG_WAITING_FOR_NOTIFY** Renamed to **PMIX_READY_FOR_DEBUG**

27 C.10.2 Deprecated attributes

28 The following attributes were deprecated in v4.2:

29 **PMIX_DEBUG_WAIT_FOR_NOTIFY** "pmix.dbg.notify" (bool)
30 Renamed to **PMIX_DEBUG_STOP_IN_APP**

1 C.10.3 Deprecated macros

2 The following macros were deprecated in v4.2:

- 3 • `PMIX_VALUE_LOAD` Replaced by the `PMIx_Value_load` API
- 4 • `PMIX_VALUE_UNLOAD` Replaced by the `PMIx_Value_unload` API
- 5 • `PMIX_VALUE_XFER` Replaced by the `PMIx_Value_xfer` API
- 6 • `PMIX_INFO_LOAD` Replaced by the `PMIx_Info_load` API
- 7 • `PMIX_INFO_XFER` Replaced by the `PMIx_Info_xfer` API
- 8 • `PMIX_INFO_LIST_START` Replaced by the `PMIx_Info_list_start` API
- 9 • `PMIX_INFO_LIST_ADD` Replaced by the `PMIx_Info_list_add` API
- 10 • `PMIX_INFO_LIST_XFER` Replaced by the `PMIx_Info_list_xfer` API
- 11 • `PMIX_INFO_LIST_CONVERT` Replaced by the `PMIx_Info_list_convert` API
- 12 • `PMIX_INFO_LIST_RELEASE` Replaced by the `PMIx_Info_list_release` API
- 13 • `PMIX_TOPOLOGY_DESTRUCT` Replaced by the `PMIx_Topology_destruct` API
- 14 • `PMIX_TOPOLOGY_FREE` Not replaced.

15 C.10.4 Added Functions (Provisional)

- 16 • `PMIx_Data_embed`
- 17 • `PMIx_Value_load`
- 18 • `PMIx_Value_unload`
- 19 • `PMIx_Value_xfer`
- 20 • `PMIx_Info_list_start`
- 21 • `PMIx_Info_list_add`
- 22 • `PMIx_Info_list_xfer`
- 23 • `PMIx_Info_list_convert`
- 24 • `PMIx_Info_list_release`
- 25 • `PMIx_Topology_destruct`

26 C.10.5 Added Macros (Provisional)

- 27 • `PMIX_APP_STATIC_INIT`
- 28 • `PMIX_BYTE_OBJECT_STATIC_INIT`
- 29 • `PMIX_COORD_STATIC_INIT`
- 30 • `PMIX_CPUSET_STATIC_INIT`
- 31 • `PMIX_DATA_ARRAY_STATIC_INIT`
- 32 • `PMIX_DATA_BUFFER_STATIC_INIT`
- 33 • `PMIX_DEVICE_DIST_STATIC_INIT`
- 34 • `PMIX_ENDPOINT_STATIC_INIT`
- 35 • `PMIX_ENVAR_STATIC_INIT`
- 36 • `PMIX_FABRIC_STATIC_INIT`

- 1 • [PMIX_GEOMETRY_STATIC_INIT](#)
- 2 • [PMIX_INFO_STATIC_INIT](#)
- 3 • [PMIX_LOOKUP_STATIC_INIT](#)
- 4 • [PMIX_PROC_INFO_STATIC_INIT](#)
- 5 • [PMIX_PROC_STATIC_INIT](#)
- 6 • [PMIX_QUERY_STATIC_INIT](#)
- 7 • [PMIX_REGATTR_STATIC_INIT](#)
- 8 • [PMIX_TOPOLOGY_STATIC_INIT](#)
- 9 • [PMIX_VALUE_STATIC_INIT](#)

10 C.10.6 Added Constants (Provisional)

11 Spawn constants

- 12 • [PMIX_ERR_JOB_EXE_NOT_FOUND](#)
- 13 • [PMIX_ERR_JOB_INSUFFICIENT_RESOURCES](#)
- 14 • [PMIX_ERR_JOB_SYS_OP_FAILED](#)
- 15 • [PMIX_ERR_JOB_WDIR_NOT_FOUND](#)

16 C.10.7 Added Attributes (Provisional)

17 Spawn attributes

18 [PMIX_ENVARS_HARVESTED](#) "pmix.evar.hvstd" (bool)

19 Environmental parameters have been harvested by the spawn requestor - the server does not
20 need to harvest them.

21 [PMIX_JOB_TIMEOUT](#) "pmix.job.time" (int)

22 Time in seconds before the spawned job should time out and be terminated (0 => infinite),
23 defined as the total runtime of the job (equivalent to the walltime limit of typical batch
24 schedulers).

25 [PMIX_LOCAL_COLLECTIVE_STATUS](#) "pmix.loc.col.st" (pmix_status_t)

26 Status code for local collective operation being reported to the host by the server library.

27 PMIx servers may aggregate the participation by local client processes in a collective
28 operation - e.g., instead of passing individual client calls to [PMIx_Fence](#) up to the host
29 environment, the server may pass only a single call to the host when all local participants
30 have executed their [PMIx_Fence](#) call, thereby reducing the burden placed on the host.

31 However, in cases where the operation locally fails (e.g., if a participating client abnormally
32 terminates prior to calling the operation), the server upcall functions to the host do not
33 include a [pmix_status_t](#) by which the PMIx server can alert the host to that failure.
34 This attribute resolves that problem by allowing the server to pass the status information
35 regarding the local collective operation.

36 [PMIX_NODE_OVERSUBSCRIBED](#) "pmix.ndosub" (bool)

1 True if the number of processes from this job on this node exceeds the number of slots
2 allocated to it

3 **PMIX_SINGLETON** "pmix.singleton" (char*)

4 String representation (nspace.rank) of proc ID for the singleton the server was started to
5 support

6 **PMIX_SPAWN_TIMEOUT** "pmix.sp.time" (int)

7 Time in seconds before spawn operation should time out (0 => infinite). Logically
8 equivalent to passing the **PMIX_TIMEOUT** attribute to the **PMIx_Spawn** API, it is
9 provided as a separate attribute to distinguish it from the **PMIX_JOB_TIMEOUT** attribute

10 Tool attributes

11 **PMIX_IOF_FILE_PATTERN** "pmix.iof.fpt" (bool)

12 Specified output file is to be treated as a pattern and not automatically annotated by nspace,
13 rank, or other parameters. The pattern can use %n for the namespace, and %r for the rank
14 wherever those quantities are to be placed. The resulting filename will be appended with
15 ".stdout" for the **stdout** stream and ".stderr" for the **stderr** stream. If

16 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
17 created and both streams will be written into it.

18 **PMIX_IOF_FILE_ONLY** "pmix.iof.fonly" (bool)

19 Output only into designated files - do not also output a copy to the console's stdout/stderr

20 **PMIX_IOF_LOCAL_OUTPUT** "pmix.iof.local" (bool)

21 Write output streams to local stdout/err

22 **PMIX_IOF_MERGE_STDERR_STDOUT** "pmix.iof.mrg" (bool)

23 Merge stdout and stderr streams from application procs

24 **PMIX_IOF_RANK_OUTPUT** "pmix.iof.rank" (bool)

25 Tag output with the rank it came from

26 **PMIX_IOF_OUTPUT_RAW** "pmix.iof.raw" (bool)

27 Do not buffer output to be written as complete lines - output characters as the stream delivers
28 them

29 **PMIX_IOF_OUTPUT_TO_DIRECTORY** "pmix.iof.dir" (char*)

30 Direct application output into files of form "<directory>/<namespace>/rank.<rank>/stdout" (for
31 **stdout**) and "<directory>/<namespace>/rank.<rank>/stderr" (for **stderr**). If

32 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
33 created and both streams will be written into it.

34 **PMIX_IOF_OUTPUT_TO_FILE** "pmix.iof.file" (char*)

35 Direct application output into files of form "<filename>.<namespace>.<rank>.stdout" (for
36 **stdout**) and "<filename>.<namespace>.<rank>.stderr" (for **stderr**). If

37 **PMIX_IOF_MERGE_STDERR_STDOUT** was given, then only the **stdout** file will be
38 created and both streams will be written into it.

1 **PMIX_BREAKPOINT** "pmix.brkpt" (char*)

2 String ID of the breakpoint where the process(es) is(are) waiting.

3 **PMIX_DEBUG_STOP_IN_APP** "pmix.dbg.notify" (varies)

4 Direct specified ranks to stop at application-specific point and notify they are
5 ready-to-debug. The attribute's value can be any of three data types:

- 6 • bool - true indicating all ranks
- 7 • **pmix_rank_t** - the rank of one proc, or **PMIX_RANK_WILDCARD** for all
- 8 • a **pmix_data_array_t** if an array of individual processes are specified

9 The resulting application processes are to notify their server (by generating the
10 **PMIX_READY_FOR_DEBUG** event) when they reach some application-determined location
11 - the event shall include the **PMIX_BREAKPOINT** attribute indicating where the application
12 has stopped. The application shall pause at that point until released by debugger
13 modification of an appropriate variable. The launcher (RM or IL) is responsible for
14 generating the **PMIX_READY_FOR_DEBUG** event when all processes have indicated they
15 are at the pause point.

APPENDIX D

Acknowledgements

1 This document represents the work of many people who have contributed to the PMIx community.
2 Without the hard work and dedication of these people this document would not have been possible.
3 The sections below list some of the active participants and organizations in the various PMIx
4 standard iterations.

5 D.1 Version 4.0

6 The following list includes some of the active participants in the PMIx v4 standardization process.

- 7 • Ralph H. Castain and Danielle Sikich
- 8 • Joshua Hursey and David Solt
- 9 • Dirk Schubert
- 10 • John DelSignore
- 11 • Aurelien Bouteiller
- 12 • Michael A Raymond
- 13 • Howard Pritchard and Nathan Hjelm
- 14 • Brice Goglin
- 15 • Kathryn Mohror and Stephen Herbein
- 16 • Thomas Naughton and Swaroop Pophale
- 17 • William E. Allcock and Paul Rich
- 18 • Michael Karo
- 19 • Artem Polyakov

20 The following institutions supported this effort through time and travel support for the people listed
21 above.

- 22 • Intel Corporation
- 23 • IBM, Inc.
- 24 • Allinea (ARM)

- 1 • Perforce
- 2 • University of Tennessee, Knoxville
- 3 • The Exascale Computing Project, an initiative of the US Department of Energy
- 4 • National Science Foundation
- 5 • HPE Co.
- 6 • Los Alamos National Laboratory
- 7 • INRIA
- 8 • Lawrence Livermore National Laboratory
- 9 • Oak Ridge National Laboratory
- 10 • Argonne National Laboratory
- 11 • Altair
- 12 • NVIDIA

13 **D.2 Version 3.0**

14 The following list includes some of the active participants in the PMIx v3 standardization process.

- 15 • Ralph H. Castain, Andrew Friedley, Brandon Yates
- 16 • Joshua Hursey and David Solt
- 17 • Aurelien Bouteiller and George Bosilca
- 18 • Dirk Schubert
- 19 • Kevin Harms
- 20 • Artem Polyakov

21 The following institutions supported this effort through time and travel support for the people listed
22 above.

- 23 • Intel Corporation
- 24 • IBM, Inc.
- 25 • University of Tennessee, Knoxville
- 26 • The Exascale Computing Project, an initiative of the US Department of Energy
- 27 • National Science Foundation
- 28 • Argonne National Laboratory

- 1 • Allinea (ARM)
- 2 • NVIDIA

3 **D.3 Version 2.0**

4 The following list includes some of the active participants in the PMIx v2 standardization process.

- 5 • Ralph H. Castain, Annapurna Dasari, Christopher A. Holguin, Andrew Friedley, Michael Klemm
- 6 and Terry Wilmarth
- 7 • Joshua Hursey, David Solt, Alexander Eichenberger, Geoff Paulsen, and Sameh Sharkawi
- 8 • Aurelien Bouteiller and George Bosilca
- 9 • Artem Polyakov, Igor Ivanov and Boris Karasev
- 10 • Gilles Gouillardet
- 11 • Michael A Raymond and Jim Stoffel
- 12 • Dirk Schubert
- 13 • Moe Jette
- 14 • Takahiro Kawashima and Shinji Sumimoto
- 15 • Howard Pritchard
- 16 • David Beer
- 17 • Brice Goglin
- 18 • Geoffroy Vallee, Swen Boehm, Thomas Naughton and David Bernholdt
- 19 • Adam Moody and Martin Schulz
- 20 • Ryan Grant and Stephen Olivier
- 21 • Michael Karo

22 The following institutions supported this effort through time and travel support for the people listed
23 above.

- 24 • Intel Corporation
- 25 • IBM, Inc.
- 26 • University of Tennessee, Knoxville
- 27 • The Exascale Computing Project, an initiative of the US Department of Energy
- 28 • National Science Foundation
- 29 • Mellanox, Inc.

- 1 • Research Organization for Information Science and Technology
- 2 • HPE Co.
- 3 • Allinea (ARM)
- 4 • SchedMD, Inc.
- 5 • Fujitsu Limited
- 6 • Los Alamos National Laboratory
- 7 • Adaptive Solutions, Inc.
- 8 • INRIA
- 9 • Oak Ridge National Laboratory
- 10 • Lawrence Livermore National Laboratory
- 11 • Sandia National Laboratory
- 12 • Altair

13 **D.4 Version 1.0**

14 The following list includes some of the active participants in the PMIx v1 standardization process.

- 15 • Ralph H. Castain, Annapurna Dasari and Christopher A. Holguin
- 16 • Joshua Hursey and David Solt
- 17 • Aurelien Bouteiller and George Bosilca
- 18 • Artem Polyakov, Elena Shipunova, Igor Ivanov, and Joshua Ladd
- 19 • Gilles Gouaillardet
- 20 • Gary Brown
- 21 • Moe Jette

22 The following institutions supported this effort through time and travel support for the people listed
23 above.

- 24 • Intel Corporation
- 25 • IBM, Inc.
- 26 • University of Tennessee, Knoxville
- 27 • Mellanox, Inc.
- 28 • Research Organization for Information Science and Technology

- 1 • Adaptive Solutions, Inc.
- 2 • SchedMD, Inc.

Unofficial Draft

Bibliography

- [1] Ralph H. Castain, David Solt, Joshua Hursey, and Aurelien Bouteiller. PMix: Process management for exascale environments. In *Proceedings of the 24th European MPI Users' Group Meeting*, EuroMPI '17, pages 14:1–14:10, New York, NY, USA, 2017. ACM.
- [2] Balaji P. et al. PMI: A scalable parallel process-management interface for extreme-scale systems. In *Recent Advances in the Message Passing Interface*, EuroMPI '10, pages 31–41, Berlin, Heidelberg, 2010. Springer.
- [3] Khaled Hamidouche, Jian Lin, Mingzhe Li, Jie Zhang, and D.K.Panda. Supporting hybrid MPI+PGAS programming models through unified communication runtime: An MVAPICH2-X approach. In *4th MVAPICH User's Group*, MUG, 2016.
- [4] Message Passing Interface Forum. *MPI: A Message-passing Interface Standard, Version 3.1*. High-Performance Computing Center Stuttgart, University of Stuttgart, June 2015.
- [5] Geoffroy R. Vallée. MOC - MPI Open MP Coordination library. <https://github.com/OMPI-X/MOC>, 2018. [Online; accessed 20-Dec-2019].
- [6] Geoffroy R. Vallée and David Bernhold. Improving support of MPI+OpenMP applications. In *Proceedings of the 25th European MPI Users' Group Meeting*, EuroMPI, 2018.

Index

General terms and other items not induced in the other indices.

application, [8](#), [96](#), [117](#), [322](#), [325](#), [327](#), [586](#), [589](#)

attribute, [9](#)

business card exchange, [72](#)

client, [5](#), [8](#), [63](#)

clients, [8](#)

clone, [8](#)

clones, [8](#), [122](#), [124](#), [125](#), [196](#), [199](#), [200](#), [202](#), [226](#), [229](#), [597](#)

data realm, [86](#), [287](#)

data realms, [86](#)

device, [9](#)

devices, [9](#)

Direct Modex, [274](#), [338](#)

fabric, [9](#)

fabric device, [9](#)

fabric devices, [9](#)

fabric plane, [9](#), [180](#), [186](#), [218](#), [221](#), [222](#), [292](#), [341](#)

fabric planes, [9](#)

fabrics, [9](#)

host environment, [8](#)

instant on, [9](#), [73](#), [273](#)

job, [7](#), [8](#), [88–92](#), [96](#), [117](#), [313–317](#), [319](#), [322](#), [324–327](#), [339](#), [341](#), [342](#), [533](#), [551](#), [586](#), [587](#), [589](#),
[600–602](#)

key, [9](#)

modex exchange, [72](#)

namespace, [7](#)

node, [8](#), [96](#), [117](#), [180](#), [186](#), [218](#), [221](#), [222](#), [322](#), [341](#)

package, [8](#), [93](#), [320](#), [602](#)

peer, [8](#), [91](#), [319](#)
peers, [8](#)
process, [8](#), [96](#), [117](#), [180](#), [186](#), [218](#), [221](#), [222](#), [322](#), [341](#)
processing unit, [9](#)

rank, [8](#), [328](#)
realm, [86](#)
realms, [86](#)
resource manager, [8](#)
RM, [8](#)

scheduler, [8](#), [290](#)
server, [5](#)
session, [7](#), [87](#), [96](#), [117](#), [313](#), [314](#), [322](#), [324](#), [586](#), [587](#), [589](#), [600](#)
singleton, [5](#)

thread, [8](#)
threads, [8](#)
tool, [5](#), [8](#), [617](#)
tools, [8](#)

workflow, [8](#)
workflows, [8](#), [411](#)

Unofficial Draft

Index of APIs

PMIx_Abort, [26](#), [174](#), [175](#), [363](#), [365](#), [488](#), [503](#), [584](#)
 PMIxClient.abort (Python), [502](#)

PMIx_Alloc_directive_string, [61](#), [523](#), [585](#)
 PMIxClient.alloc_directive_string (Python), [523](#)

PMIx_Allocation_request, [119](#), [215](#), [215](#), [221](#), [511](#), [587](#), [589](#), [592](#)
 PMIxClient.allocation_request (Python), [510](#)

PMIx_Allocation_request_nb, [218](#), [221](#), [223](#), [585](#)

PMIx_Commit, [73](#), [75](#), [76](#), [77](#), [78](#), [121](#), [123](#), [124](#), [338](#), [339](#), [366](#), [370](#), [504](#), [549](#), [550](#), [584](#)
 PMIxClient.commit (Python), [504](#)

PMIx_Compute_distances, [208](#), [210](#), [520](#), [593](#)
 PMIxClient.compute_distances (Python), [520](#)

PMIx_Compute_distances_nb, [209](#), [593](#)

PMIx_Connect, [195](#), [197](#), [199](#), [200](#), [202](#), [247–249](#), [457](#), [458](#), [508](#), [584](#), [586](#)
 PMIxClient.connect (Python), [507](#)

PMIx_Connect_nb, [197](#), [197](#), [584](#)

pmix_connection_cbfunc_t, [392](#), [392](#)

pmix_credential_cbfunc_t, [298](#), [409](#), [410](#)

PMIx_Data_compress, [171](#), [172](#), [614](#)

PMIx_Data_copy, [167](#), [585](#)

PMIx_Data_copy_payload, [168](#), [585](#)

PMIx_Data_decompress, [172](#), [614](#)

PMIx_Data_embed, [173](#), [618](#)

PMIx_Data_load, [169](#), [173](#), [614](#)

PMIx_Data_pack, [164](#), [165](#), [311](#), [585](#)

PMIx_Data_print, [168](#), [585](#)

PMIx_Data_range_string, [60](#), [522](#), [585](#)
 PMIxClient.data_range_string (Python), [522](#)

PMIx_Data_type_string, [60](#), [523](#), [585](#)
 PMIxClient.data_type_string (Python), [523](#)

PMIx_Data_unload, [170](#), [614](#)

PMIx_Data_unpack, [165](#), [171](#), [585](#)

PMIx_Deregister_event_handler, [155](#), [516](#), [568](#), [585](#), [592](#)
 PMIxClient.deregister_event_handler (Python), [516](#)

pmix_device_dist_cbfunc_t, [210](#), [210](#), [594](#)

PMIx_Device_type_string, [62](#), [526](#), [593](#)
 PMIxClient.device_type_string (Python), [526](#)

PMIx_Disconnect, [199](#), [200–202](#), [249](#), [457](#), [458](#), [508](#), [584](#), [586](#)
 PMIxClient.disconnect (Python), [508](#)

PMIx_Disconnect_nb, [201](#), [202](#), [249](#), [584](#)
pmix_dmodex_response_fn_t, [338](#), [339](#)
PMIx_Error_string, [59](#), [521](#), [584](#)
 PMIxClient.error_string (Python), [520](#)
pmix_event_notification_cbfunc_fn_t, [154](#), [160](#), [160](#), [573](#), [576](#)
PMIx_Fabric_deregister, [294](#), [295](#), [518](#), [593](#)
 PMIxClient.fabric_deregister (Python), [518](#)
PMIx_Fabric_deregister_nb, [295](#), [593](#)
PMIx_Fabric_register, [283](#), [290](#), [293](#), [517](#), [593](#)
 PMIxClient.fabric_register (Python), [517](#)
PMIx_Fabric_register_nb, [292](#), [593](#)
PMIx_Fabric_update, [292](#), [293](#), [294](#), [518](#), [593](#)
 PMIxClient.fabric_update (Python), [517](#)
PMIx_Fabric_update_nb, [294](#), [593](#)
PMIx_Fence, [4](#), [73](#), [120](#), [120–122](#), [124](#), [125](#), [197](#), [200](#), [246](#), [255](#), [259](#), [274](#), [308](#), [338](#), [365–367](#), [383](#),
 [386](#), [421](#), [489](#), [505](#), [549](#), [550](#), [584](#), [597](#), [619](#)
 PMIxClient.fence (Python), [504](#)
PMIx_Fence_nb, [57](#), [120](#), [122](#), [365](#), [367](#), [489](#), [584](#), [592](#)
PMIx_Finalize, [26](#), [67](#), [69](#), [69](#), [195](#), [362](#), [363](#), [457](#), [458](#), [488](#), [502](#), [584](#)
 PMIxClient.finalize (Python), [502](#)
PMIx_generate_ppn, [311](#), [528](#), [584](#), [590](#)
 PMIxServer.generate_ppn (Python), [528](#)
PMIx_generate_regex, [310](#), [312](#), [324](#), [528](#), [584](#), [590](#)
 PMIxServer.generate_regex (Python), [527](#)
PMIx_Get, [3](#), [9](#), [29](#), [67](#), [72–74](#), [77](#), [78](#), [80](#), [81](#), [83–89](#), [91–96](#), [100](#), [103](#), [106–108](#), [112](#), [117](#),
 [119–121](#), [124](#), [125](#), [177–179](#), [183–185](#), [188–190](#), [204](#), [206](#), [207](#), [221](#), [230](#), [236](#), [243](#), [245](#),
 [247](#), [251](#), [255](#), [274](#), [276](#), [287](#), [290](#), [315](#), [316](#), [320](#), [323](#), [353](#), [379](#), [381](#), [431](#), [441–443](#), [459](#),
 [476](#), [505](#), [549](#), [550](#), [552](#), [555](#), [557](#), [561](#), [562](#), [567](#), [579](#), [584](#), [586](#), [591](#), [597](#), [598](#), [604](#), [605](#),
 [614](#)
 PMIxClient.get (Python), [505](#)
PMIx_Get_attribute_name, [61](#), [525](#), [593](#)
 PMIxClient.get_attribute_name (Python), [525](#)
PMIx_Get_attribute_string, [61](#), [525](#), [593](#)
 PMIxClient.get_attribute_string (Python), [524](#)
PMIx_Get_cpuset, [207](#), [519](#), [593](#)
 PMIxClient.get_cpuset (Python), [519](#)
PMIx_Get_credential, [297](#), [299](#), [410](#), [512](#), [587](#), [592](#)
 PMIxClient.get_credential (Python), [512](#)
PMIx_Get_credential_nb, [298](#)
PMIx_Get_nb, [57](#), [58](#), [72](#), [73](#), [81](#), [120](#), [584](#)
PMIx_Get_relative_locality, [204](#), [206](#), [320](#), [354](#), [519](#), [552](#), [593](#), [613](#)
 PMIxClient.get_relative_locality (Python), [519](#)
PMIx_Get_version, [11](#), [64](#), [502](#), [584](#)

PMIxCliet.get_version (Python), 502
 PMIx_Group_construct, 246, 247, [252](#), 254, 255, 258, 513, 579, 593
 PMIxCliet.group_construct (Python), 513
 PMIx_Group_construct_nb, [255](#), 258, 593
 PMIx_Group_destruct, 249, [258](#), 259, 261, 271, 515, 593
 PMIxCliet.group_destruct (Python), 515
 PMIx_Group_destruct_nb, [259](#), 261, 593
 PMIx_Group_invite, 248, [261](#), 263, 264, 266, 514, 593
 PMIxCliet.group_invite (Python), 513
 PMIx_Group_invite_nb, [264](#), 593
 PMIx_Group_join, 248, 263, [266](#), 266–268, 270, 514, 593
 PMIxCliet.group_join (Python), 514
 PMIx_Group_join_nb, 266, [268](#), 270, 593
 PMIx_Group_leave, 249, [270](#), 271, 272, 515, 593
 PMIxCliet.group_leave (Python), 515
 PMIx_Group_leave_nb, [271](#), 593
 pmix_hdlr_reg_cbfunc_t, [58](#), 149, 469, 470, 611
 pmix_info_cbfunc_t, [56](#), [58](#), 58, 106, 209, 219, 227, 229, 234, 235, [256](#), 264, 269, 351, 393, 401, 403, 404, 406, 407, 420, 423
 PMIx_Info_directives_string, [60](#), 523, 585
 PMIxCliet.info_directives_string (Python), [522](#)
 PMIx_Info_list_add, [38](#), 618
 PMIx_Info_list_convert, [39](#), 618
 PMIx_Info_list_release, [39](#), 618
 PMIx_Info_list_start, [37](#), 38–40, 618
 PMIx_Info_list_xfer, [38](#), 618
 PMIx_Info_load, [36](#), 618
 PMIx_Info_xfer, [36](#), 324, 618
 PMIx_Init, 8, 63, [64](#), 67, 93, 103, 108, 362, 436, 440, 451, 452, 459, 501, 549, 551, 556, 562, 571, 581, 585, 591, 614
 PMIxCliet.init (Python), 501
 PMIx_Initialized, [63](#), 501, 584
 PMIxCliet.initialized (Python), 501
 pmix_iof_cbfunc_t, [417](#), 468, 486
 iofcbfunc (Python), 486
 PMIx_IOF_channel_string, [61](#), 524, 587
 PMIxCliet.iof_channel_string (Python), 524
 PMIx_IOF_deregister, [470](#), 541, 587, 592
 PMIxCliet.iof_deregister (Python), 541
 PMIx_IOF_pull, 380, 397, 433, 434, 439, 443, 445, [468](#), 470, 541, 557, 561, 562, 567, 587, 592
 PMIxCliet.iof_pull (Python), 540
 PMIx_IOF_push, 380, 397, 433, 439, 442, 445, 447, 450, [471](#), 473, 542, 587, 592, 603
 PMIxCliet.iof_push (Python), 541

PMIx_Job_control, [215](#), [224](#), [225](#), [228–230](#), [405](#), [454](#), [511](#), [587](#), [592](#)
 PMIxClient.job_ctrl (Python), [511](#)
 PMIx_Job_control_nb, [100](#), [223](#), [226](#), [322](#), [585](#)
 PMIx_Job_state_string, [61](#), [524](#), [593](#)
 PMIxClient.job_state_string (Python), [524](#)
 PMIx_Link_state_string, [62](#), [525](#), [593](#)
 PMIxClient.link_state_string (Python), [525](#)
 PMIx_Load_topology, [203](#), [518](#), [593](#), [613](#)
 PMIxClient.load_topology (Python), [518](#)
 PMIx_Log, [191](#), [237](#), [239](#), [243](#), [436](#), [456](#), [457](#), [510](#), [587](#), [610](#)
 PMIxClient.log (Python), [510](#)
 PMIx_Log_nb, [240](#), [243](#), [585](#)
 PMIx_Lookup, [72](#), [126](#), [130](#), [135](#), [137](#), [139](#), [505](#), [506](#), [584](#)
 PMIxClient.lookup (Python), [506](#)
 pmix_lookup_cbfunc_t, [56](#), [56](#), [373](#)
 PMIx_Lookup_nb, [56](#), [57](#), [72](#), [137](#), [584](#)
 pmix_modex_cbfunc_t, [56](#), [365](#), [368](#), [368](#), [369](#)
 pmix_notification_fn_t, [149](#), [154](#), [154](#), [487](#)
 evhandler (Python), [486](#)
 PMIx_Notify_event, [134](#), [156](#), [391](#), [517](#), [555](#), [560](#), [568](#), [572](#), [573](#), [585](#), [592](#)
 PMIxClient.notify_event (Python), [516](#)
 pmix_op_cbfunc_t, [57](#), [57](#), [128](#), [131](#), [144](#), [156](#), [157](#), [160](#), [198](#), [201](#), [240](#), [260](#), [271](#), [293–295](#), [312](#),
 [332–336](#), [343](#), [349](#), [350](#), [352](#), [360–362](#), [364](#), [371](#), [375](#), [382](#), [385](#), [387](#), [389](#), [390](#), [399](#), [415](#),
 [418](#), [471](#)
 PMIx_Parse_cpuset_string, [206](#), [354](#), [520](#), [593](#)
 PMIxClient.parse_cpuset_string (Python), [519](#)
 PMIx_Persistence_string, [60](#), [522](#), [585](#)
 PMIxClient.persistence_string (Python), [522](#)
 PMIx_Proc_state_string, [59](#), [521](#), [585](#)
 PMIxClient.proc_state_string (Python), [521](#)
 PMIx_Process_monitor, [215](#), [231](#), [235](#), [512](#), [587](#), [592](#)
 PMIxClient.monitor (Python), [511](#)
 PMIx_Process_monitor_nb, [233](#), [236](#), [585](#)
 PMIx_Progress, [66](#), [69](#), [307](#), [310](#), [463](#), [526](#), [593](#), [602](#)
 PMIxClient.progress (Python), [526](#)
 PMIx_Publish, [72](#), [130](#), [131](#), [133–135](#), [137](#), [372](#), [506](#), [584](#)
 PMIxClient.publish (Python), [505](#)
 PMIx_Publish2, [126](#), [128](#), [130](#)
 PMIx_Publish2_nb, [128](#)
 PMIx_Publish_nb, [131](#), [135](#), [137](#), [584](#)
 PMIx_Put, [29](#), [72](#), [73](#), [74](#), [74–76](#), [78](#), [117](#), [120–124](#), [195](#), [255](#), [263](#), [338](#), [339](#), [366](#), [370](#), [504](#), [549](#),
 [550](#), [584](#)
 PMIxClient.put (Python), [503](#)

PMIx_Query_info, [9](#), [101](#), [101](#), [106](#), [110](#), [112](#), [115](#), [117](#), [118](#), [245](#), [247](#), [287](#), [425](#), [429](#), [452](#), [453](#), [476](#), [510](#), [555](#), [561](#), [564](#), [581](#), [582](#), [598](#)
 PMIxClient.query (Python), [509](#)
PMIx_Query_info_nb, [100](#), [106](#), [106](#), [110](#), [112](#), [115](#), [119](#), [195](#), [323](#), [344](#), [567](#), [585](#), [586](#)
PMIx_Register_attributes, [344](#), [534](#), [592](#), [593](#)
 PMIxServer.register_attributes (Python), [533](#)
PMIx_Register_event_handler, [100](#), [148](#), [190](#), [435](#), [457](#), [516](#), [555](#), [560](#), [564](#), [568](#), [572](#), [573](#), [576](#), [585](#), [592](#), [609](#)
 PMIxClient.register_event_handler (Python), [515](#)
pmix_release_cbfunc_t, [56](#), [56](#)
PMIx_Resolve_nodes, [116](#), [509](#), [584](#)
 PMIxClient.resolve_nodes (Python), [509](#)
PMIx_Resolve_peers, [90](#), [95](#), [115](#), [318](#), [319](#), [509](#), [551](#), [584](#)
 PMIxClient.resolve_peers (Python), [508](#)
PMIx_Scope_string, [60](#), [521](#), [585](#)
 PMIxClient.scope_string (Python), [521](#)
pmix_server_abort_fn_t, [363](#), [489](#)
 clientaborted (Python), [488](#)
pmix_server_alloc_fn_t, [400](#), [496](#)
 allocate (Python), [496](#)
pmix_server_client_connected2_fn_t, [57](#), [296](#), [336](#), [360](#), [361](#), [362](#), [488](#), [581](#), [582](#), [593](#), [611](#)
 clientconnected2 (Python), [487](#)
pmix_server_client_finalized_fn_t, [362](#), [363](#), [488](#)
 clientfinalized (Python), [488](#)
PMIx_server_collect_inventory, [351](#), [353](#), [535](#), [587](#)
 PMIxServer.collect_inventory (Python), [535](#)
pmix_server_connect_fn_t, [195](#), [382](#), [384](#), [386](#), [492](#)
 connect (Python), [492](#)
PMIx_server_define_process_set, [245](#), [356](#), [536](#), [593](#)
 PMIxServer.define_process_set (Python), [536](#)
PMIx_server_delete_process_set, [245](#), [357](#), [537](#), [593](#)
 PMIxServer.delete_process_set (Python), [536](#)
PMIx_server_deliver_inventory, [352](#), [536](#), [587](#)
 PMIxServer.deliver_inventory (Python), [535](#)
PMIx_server_deregister_client, [336](#), [532](#), [584](#)
 PMIxServer.deregister_client (Python), [532](#)
pmix_server_deregister_events_fn_t, [388](#), [494](#)
 deregister_events (Python), [494](#)
PMIx_server_deregister_nspace, [331](#), [337](#), [530](#), [584](#)
 PMIxServer.deregister_nspace (Python), [530](#)
PMIx_server_deregister_resources, [334](#), [531](#), [537](#), [594](#)
 PMIxServer.deregister_resources (Python), [531](#), [537](#)
pmix_server_disconnect_fn_t, [384](#), [386](#), [493](#)

disconnect (Python), 493

pmix_server_dmodex_req_fn_t, 78, 95, [368](#), 368, 490, 587, 589
dmodex (Python), 490

PMIx_server_dmodex_request, [337](#), 338, 339, 533, 584
PMIxServer.dmodex_request (Python), 532

pmix_server_fabric_fn_t, 283, 290, [422](#), 500, 593
fabric (Python), 500

pmix_server_fencenb_fn_t, [365](#), 367, 368, 489, 589
fence (Python), 489

PMIx_server_finalize, [308](#), 527, 584
PMIxServer.finalize (Python), 527

PMIx_server_generate_cpuset_string, 207, [354](#), 529, 594
PMIxServer.generate_cpuset_string (Python), 529

PMIx_server_generate_locality_string, 203, 204, [353](#), 529, 593
PMIxServer.generate_locality_string (Python), 528

pmix_server_get_cred_fn_t, [409](#), 413, 498
get_credential (Python), 497

pmix_server_grp_fn_t, [419](#), 500, 593
group (Python), 499

PMIx_server_init, 63, [304](#), 309, 344, 358, 426, 427, 431, 527, 568, 581, 582, 584, 593
PMIxServer.init (Python), 526

PMIx_server_IOF_deliver, [350](#), 444, 535, 587
PMIxServer.iof_deliver (Python), 534

pmix_server_iof_fn_t, [414](#), 499
iof_pull (Python), 498

pmix_server_job_control_fn_t, [403](#), 497
job_control (Python), 496

pmix_server_listener_fn_t, [391](#)

pmix_server_log_fn_t, [398](#), 496
log (Python), 495

pmix_server_lookup_fn_t, [372](#), 491
lookup (Python), 491

pmix_server_module_t, 305, 308, 344, 345, [358](#), 358, 359, 527

pmix_server_monitor_fn_t, [406](#), 497
monitor (Python), 497

pmix_server_notify_event_fn_t, 155, 159, [390](#), 391, 494
notify_event (Python), 494

pmix_server_publish_fn_t, [370](#), 490
publish (Python), 490

pmix_server_query_fn_t, [393](#), 495
query (Python), 494

PMIx_server_register_client, 296, [335](#), 336, 362, 363, 531, 584
PMIxServer.register_client (Python), 531

pmix_server_register_events_fn_t, [387](#), [493](#)
register_events (Python), [493](#)

PMIx_server_register_namespace, [11](#), [57](#), [311](#), [312](#), [313](#), [322](#), [323](#), [326](#), [333](#), [350](#), [354](#), [530](#), [584](#), [586](#), [600](#)
PMIxServer.register_namespace (Python), [529](#)

PMIx_server_register_resources, [314](#), [315](#), [318](#), [319](#), [333](#), [531](#), [537](#), [594](#)
PMIxServer.register_resources (Python), [530](#), [537](#)

PMIx_server_setup_application, [339](#), [343](#), [349](#), [353](#), [533](#), [585](#), [589](#)
PMIxServer.setup_application (Python), [533](#)

PMIx_server_setup_fork, [337](#), [532](#), [584](#)
PMIxServer.setup_fork (Python), [532](#)

PMIx_server_setup_local_support, [348](#), [534](#), [585](#)
PMIxServer.setup_local_support (Python), [534](#)

pmix_server_spawn_fn_t, [194](#), [377](#), [437](#), [492](#)
spawn (Python), [492](#)

pmix_server_stdin_fn_t, [418](#), [499](#)
iof_push (Python), [499](#)

pmix_server_tool_connection_fn_t, [296](#), [395](#), [426](#), [495](#), [581](#), [582](#)
tool_connected (Python), [495](#)

pmix_server_unpublish_fn_t, [375](#), [491](#)
unpublish (Python), [491](#)

pmix_server_validate_cred_fn_t, [411](#), [498](#)
validate_credential (Python), [498](#)

pmix_setup_application_cbfunc_t, [340](#), [343](#)

PMIx_Spawn, [89](#), [90](#), [93](#), [94](#), [175](#), [176](#), [180](#), [181](#), [186–188](#), [191](#), [192](#), [218](#), [221](#), [320](#), [322](#), [337](#), [377](#), [378](#), [380–382](#), [402](#), [431](#), [433](#), [437–439](#), [441–443](#), [451–455](#), [459](#), [492](#), [507](#), [553](#), [555](#), [557](#), [560](#), [563](#), [564](#), [566](#), [569](#), [584](#), [590](#), [601](#), [617](#), [620](#)
PMIxClient.spawn (Python), [507](#)

pmix_spawn_cbfunc_t, [182](#), [194](#), [194](#), [377](#)

PMIx_Spawn_nb, [93](#), [181](#), [192](#), [194](#), [378](#), [584](#)

PMIx_Store_internal, [74](#), [76](#), [503](#), [584](#)
PMIxClient.store_internal (Python), [503](#)

PMIx_tool_attach_to_server, [428](#), [431](#), [441](#), [463](#), [465](#), [539](#), [560](#), [581](#), [594](#), [611](#)
PMIxTool.attach_to_server (Python), [539](#)

PMIx_tool_connect_to_server, [587](#), [612](#)

pmix_tool_connection_cbfunc_t, [396](#), [397](#), [398](#)

PMIx_tool_disconnect, [464](#), [539](#), [594](#)
PMIxTool.disconnect (Python), [538](#)

PMIx_tool_finalize, [463](#), [538](#), [585](#)
PMIxTool.finalize (Python), [538](#)

PMIx_tool_get_servers, [466](#), [540](#), [594](#)
PMIxTool.get_servers (Python), [539](#)

PMIx_tool_init, [8](#), [63](#), [93](#), [103](#), [108](#), [425](#), [428](#), [430–432](#), [441](#), [442](#), [444](#), [458](#), [461](#), [463](#), [538](#), [553](#),

560, 564, 581, 585
PMIxTool.init (Python), 538
PMIx_tool_set_server, 427, 441, 442, 466, [467](#), 540, 594
PMIxTool.set_server (Python), 540
PMIx_Topology_destruct, [205](#), 618
PMIx_Unpublish, [142](#), 143, 145, 507, 584
PMIxClient.unpublish (Python), 506
PMIx_Unpublish_nb, [143](#), 584
PMIx_Validate_credential, [300](#), 513, 587, 592
PMIxClient.validate_credential (Python), 512
PMIx_Validate_credential_nb, [301](#)
pmix_validation_cbfunc_t, 302, 412, [413](#)
pmix_value_cbfunc_t, [57](#), 57
PMIx_Value_load, [32](#), 618
PMIx_Value_unload, [32](#), 618
PMIx_Value_xfer, [33](#), 618

pmix_evhdr_reg_cbfunc_t
(**Deprecated**), [611](#)
pmix_server_client_connected_fn_t
(**Deprecated**), [360](#), 611
PMIx_tool_connect_to_server
(**Deprecated**), [611](#)

Index of Support Macros

[PMIX_APP_CONSTRUCT](#), [192](#)
[PMIX_APP_CREATE](#), [193](#)
[PMIX_APP_DESTRUCT](#), [193](#)
[PMIX_APP_FREE](#), [193](#)
[PMIX_APP_INFO_CREATE](#), [194](#), [587](#), [589](#)
[PMIX_APP_RELEASE](#), [193](#)
[PMIX_APP_STATIC_INIT](#), [192](#), [618](#)
[PMIX_ARGV_APPEND](#), [49](#)
[PMIX_ARGV_APPEND_UNIQUE](#), [50](#)
[PMIX_ARGV_COPY](#), [52](#)
[PMIX_ARGV_COUNT](#), [52](#)
[PMIX_ARGV_FREE](#), [51](#)
[PMIX_ARGV_JOIN](#), [52](#)
[PMIX_ARGV_PREPEND](#), [50](#)
[PMIX_ARGV_SPLIT](#), [51](#)
[PMIX_BYTE_OBJECT_CONSTRUCT](#), [46](#)
[PMIX_BYTE_OBJECT_CREATE](#), [46](#)
[PMIX_BYTE_OBJECT_DESTRUCT](#), [46](#)
[PMIX_BYTE_OBJECT_FREE](#), [47](#)
[PMIX_BYTE_OBJECT_LOAD](#), [47](#)
[PMIX_BYTE_OBJECT_STATIC_INIT](#), [46](#), [618](#)
[PMIX_CHECK_KEY](#), [18](#)
[PMIX_CHECK_NAMESPACE](#), [19](#)
[PMIX_CHECK_PROCID](#), [24](#)
[PMIX_CHECK_RANK](#), [21](#)
[PMIX_CHECK_RESERVED_KEY](#), [18](#), [610](#)
[PMIX_COORD_CONSTRUCT](#), [279](#)
[PMIX_COORD_CREATE](#), [279](#)
[PMIX_COORD_DESTRUCT](#), [279](#)
[PMIX_COORD_FREE](#), [280](#)
[PMIX_COORD_STATIC_INIT](#), [279](#), [618](#)
[PMIX_CPUSSET_CONSTRUCT](#), [355](#)
[PMIX_CPUSSET_CREATE](#), [356](#)
[PMIX_CPUSSET_DESTRUCT](#), [355](#)
[PMIX_CPUSSET_FREE](#), [356](#)
[PMIX_CPUSSET_STATIC_INIT](#), [355](#), [618](#)
[PMIX_DATA_ARRAY_CONSTRUCT](#), [48](#)
[PMIX_DATA_ARRAY_CREATE](#), [48](#)

PMIX_DATA_ARRAY_DESTRUCT, [48](#)
PMIX_DATA_ARRAY_FREE, [49](#)
PMIX_DATA_ARRAY_STATIC_INIT, [48](#), [618](#)
PMIX_DATA_BUFFER_CONSTRUCT, [162](#), [165](#), [166](#)
PMIX_DATA_BUFFER_CREATE, [162](#), [165](#), [166](#)
PMIX_DATA_BUFFER_DESTRUCT, [163](#)
PMIX_DATA_BUFFER_LOAD, [163](#)
PMIX_DATA_BUFFER_RELEASE, [162](#)
PMIX_DATA_BUFFER_STATIC_INIT, [162](#), [618](#)
PMIX_DATA_BUFFER_UNLOAD, [163](#), [311](#)
PMIX_DEVICE_DIST_CONSTRUCT, [212](#)
PMIX_DEVICE_DIST_CREATE, [213](#)
PMIX_DEVICE_DIST_DESTRUCT, [213](#)
PMIX_DEVICE_DIST_FREE, [213](#)
PMIX_DEVICE_DIST_STATIC_INIT, [212](#), [618](#)
PMIX_ENDPOINT_CONSTRUCT, [277](#)
PMIX_ENDPOINT_CREATE, [278](#)
PMIX_ENDPOINT_DESTRUCT, [277](#)
PMIX_ENDPOINT_FREE, [278](#)
PMIX_ENDPOINT_STATIC_INIT, [277](#), [618](#)
PMIX_ENVAR_CONSTRUCT, [44](#)
PMIX_ENVAR_CREATE, [44](#)
PMIX_ENVAR_DESTRUCT, [13](#), [44](#)
PMIX_ENVAR_FREE, [45](#)
PMIX_ENVAR_LOAD, [45](#)
PMIX_ENVAR_STATIC_INIT, [44](#), [618](#)
PMIX_FABRIC_CONSTRUCT, [286](#)
PMIX_FABRIC_STATIC_INIT, [286](#), [618](#)
PMIX_GEOMETRY_CONSTRUCT, [281](#)
PMIX_GEOMETRY_CREATE, [281](#)
PMIX_GEOMETRY_DESTRUCT, [281](#)
PMIX_GEOMETRY_FREE, [282](#)
PMIX_GEOMETRY_STATIC_INIT, [281](#), [619](#)
PMIx_Heartbeat, [235](#), [585](#)
PMIX_INFO_CONSTRUCT, [35](#)
PMIX_INFO_CREATE, [35](#), [41](#), [43](#)
PMIX_INFO_DESTRUCT, [35](#)
PMIX_INFO_FREE, [36](#)
PMIX_INFO_IS_END, [43](#), [587](#), [589](#)
PMIX_INFO_IS_OPTIONAL, [42](#)
PMIX_INFO_IS_REQUIRED, [40](#), [41](#), [42](#)
PMIX_INFO_LIST_ADD, [610](#)
PMIX_INFO_LIST_CONVERT, [610](#)

PMIX_INFO_LIST_RELEASE, [610](#)
PMIX_INFO_LIST_START, [610](#)
PMIX_INFO_LIST_XFER, [610](#)
PMIX_INFO_OPTIONAL, [42](#)
PMIX_INFO_PROCESSED, [42](#), [610](#)
PMIX_INFO_REQUIRED, [40](#), [41](#)
PMIX_INFO_STATIC_INIT, [35](#), [619](#)
PMIX_INFO_TRUE, [37](#)
PMIX_INFO_WAS_PROCESSED, [43](#), [610](#)
PMIX_LOAD_KEY, [18](#)
PMIX_LOAD_NAMESPACE, [20](#)
PMIX_LOAD_PROCID, [23](#), [24](#)
PMIX_LOOKUP_STATIC_INIT, [139](#), [619](#)
PMIX_MULTICLUSTER_NAMESPACE_CONSTRUCT, [25](#)
PMIX_MULTICLUSTER_NAMESPACE_PARSE, [25](#)
PMIX_NAMESPACE_INVALID, [20](#), [615](#)
PMIX_PDATA_CONSTRUCT, [139](#)
PMIX_PDATA_CREATE, [139](#)
PMIX_PDATA_DESTRUCT, [139](#)
PMIX_PDATA_FREE, [140](#)
PMIX_PDATA_LOAD, [140](#)
PMIX_PDATA_RELEASE, [140](#)
PMIX_PDATA_XFER, [141](#)
PMIX_PROC_CONSTRUCT, [22](#)
PMIX_PROC_CREATE, [22](#)
PMIX_PROC_DESTRUCT, [22](#)
PMIX_PROC_FREE, [23](#), [116](#)
PMIX_PROC_INFO_CONSTRUCT, [27](#)
PMIX_PROC_INFO_CREATE, [28](#)
PMIX_PROC_INFO_DESTRUCT, [28](#)
PMIX_PROC_INFO_FREE, [28](#)
PMIX_PROC_INFO_RELEASE, [28](#)
PMIX_PROC_INFO_STATIC_INIT, [27](#), [619](#)
PMIX_PROC_LOAD, [23](#)
PMIX_PROC_RELEASE, [23](#)
PMIX_PROC_STATIC_INIT, [22](#), [619](#)
PMIX_PROCID_INVALID, [24](#), [615](#)
PMIX_PROCID_XFER, [25](#), [615](#)
PMIX_QUERY_CONSTRUCT, [114](#)
PMIX_QUERY_CREATE, [114](#)
PMIX_QUERY_DESTRUCT, [114](#)
PMIX_QUERY_FREE, [115](#)
PMIX_QUERY_QUALIFIERS_CREATE, [115](#), [587](#), [589](#)

PMIX_QUERY_RELEASE, [114](#)
PMIX_QUERY_STATIC_INIT, [113](#), [619](#)
PMIX_RANK_IS_VALID, [21](#), [615](#)
PMIX_REGATTR_CONSTRUCT, [347](#)
PMIX_REGATTR_CREATE, [347](#)
PMIX_REGATTR_DESTRUCT, [347](#)
PMIX_REGATTR_FREE, [347](#)
PMIX_REGATTR_LOAD, [348](#)
PMIX_REGATTR_STATIC_INIT, [346](#), [619](#)
PMIX_REGATTR_XFER, [348](#)
PMIX_SETENV, [53](#)
PMIX_SYSTEM_EVENT, [152](#)
PMIX_TOPOLOGY_CONSTRUCT, [205](#)
PMIX_TOPOLOGY_CREATE, [205](#)
PMIX_TOPOLOGY_STATIC_INIT, [205](#), [619](#)
PMIX_VALUE_CONSTRUCT, [31](#)
PMIX_VALUE_CREATE, [31](#)
PMIX_VALUE_DESTRUCT, [31](#), [80](#), [84](#), [597](#)
PMIX_VALUE_FREE, [32](#)
PMIX_VALUE_GET_NUMBER, [34](#)
PMIX_VALUE_RELEASE, [31](#)
PMIX_VALUE_STATIC_INIT, [31](#), [619](#)

PMIX_INFO_LIST_ADD
(Deprecated), [618](#)
PMIX_INFO_LIST_CONVERT
(Deprecated), [618](#)
PMIX_INFO_LIST_RELEASE
(Deprecated), [618](#)
PMIX_INFO_LIST_START
(Deprecated), [618](#)
PMIX_INFO_LIST_XFER
(Deprecated), [618](#)
PMIX_INFO_LOAD
(Deprecated), [618](#)
PMIX_INFO_XFER
(Deprecated), [618](#)
PMIX_TOPOLOGY_DESTRUCT
(Deprecated), [618](#)
PMIX_TOPOLOGY_FREE
(Deprecated), [618](#)
PMIX_VALUE_LOAD
(Deprecated), [618](#)
PMIX_VALUE_UNLOAD

(Deprecated), [618](#)
PMIX_VALUE_XFER
(Deprecated), [618](#)

Unofficial Draft

Index of Data Structures

pmix_alloc_directive_t, 55, 61, 216, 219, **223**, 223, 401, 482, 523
pmix_app_t, 49, 50, 53, 176, 178, 182, 184, 189, **192**, 192–194, 377, 378, 380, 433, 436, 438, 440, 442, 451, 454, 455, 459, 460, 483, 556, 557, 562, 563, 566, 587, 589, 609
pmix_bind_envelope_t, 207, **208**, 208, 484, 594
pmix_byte_object_t, **45**, 45–47, 54, 169–171, 173, 297, 298, 300, 302, 350, 411, 412, 418, 471, 481
pmix_coord_t, 55, **278**, 278–280, 484, 594
pmix_coord_view_t, **282**, 484, 594
pmix_cpuset_t, 55, 208, 209, 353, 354, **355**, 355, 356, 483, 594
pmix_data_array_t, 29, 39, **47**, 47–49, 54, 91, 104, 106, 109, 111, 112, 118, 217, 220, 222, 247, 250, 284, 285, 287, 288, 290, 291, 314–316, 318, 319, 326–328, 341, 394, 403, 424, 436, 452, 453, 459, 460, 482, 558, 563, 566, 579, 587, 589, 605–608, 621
pmix_data_buffer_t, **161**, 161–165, 169, 170
pmix_data_range_t, 54, 60, **134**, 134, 157, 390, 481, 522
pmix_data_type_t, 32, 34, 36, 38, 48, **53**, 53, 54, 60, 140, 164, 166–168, 348, 480, 523
pmix_device_distance_t, 55, 208, 210, **211**, 211–214, 321, 484, 594, 607
pmix_device_type_t, 55, 62, **211**, 211, 290, 484, 526, 594
pmix_endpoint_t, 55, **276**, 276–278, 290, 483, 605
pmix_envvar_t, 13, **43**, 43–45, 55, 482, 569
pmix_fabric_operation_t, **283**, 283, 423
pmix_fabric_t, 276, **283**, 283, 286, 288, 291–295, 424, 483, 594, 604
pmix_geometry_t, 55, 275, **280**, 280–282, 288, 289, 484, 594, 605, 606
pmix_group_operation_t, 420, **422**, 422, 594
pmix_group_opt_t, 266, 269, **270**, 270, 514, 594
pmix_info_directives_t, **40**, 40, 41, 54, 60, 482, 522
pmix_info_t, 4, 5, 9, 18, **34**, 34–43, 54, 58, 64, 67, 69, 101, 106, 112, 113, 115, 118, 126–129, 131–135, 154, 157, 160, 194, 208, 209, 216, 217, 219, 220, 222–224, 226, 229, 230, 232, 235, 239, 242, 244, 252, 254, 256, 258, 260–262, 264, 266, 269–271, 284, 285, 288, 290, 297, 298, 300, 302, 305, 308, 313–315, 317, 318, 322–324, 326–328, 340, 341, 346, 348, 350–352, 358, 361, 380, 390, 396, 397, 400, 402, 403, 406, 407, 414, 415, 417, 424, 433, 436, 438, 440, 443, 451, 459, 461, 465, 467, 468, 470, 471, 482, 485, 556, 557, 561–563, 566, 574, 585, 587, 589, 598–601, 604–606
pmix_iof_channel_t, 55, 61, 350, 415, 417, **448**, 448, 468, 482, 524
pmix_job_state_t, **29**, 29, 55, 61, 483, 524, 594
pmix_key_t, 9, **17**, 17, 74, 79, 348, 480
pmix_link_state_t, 55, 62, 276, **282**, 282, 283, 286, 289, 483, 525, 594, 607
pmix_locality_t, 55, 204, **206**, 206, 483, 594, 613
pmix_nspace_t, **19**, 19, 20, 23–25, 54, 194, 480, 481
pmix_pdata_t, 56, 57, **135**, 135, 137, 139–141, 482

pmix_persistence_t, 54, 60, [134](#), 134, 481, 522
pmix_proc_info_t, [27](#), 27, 28, 54, 103, 104, 108, 109, 111, 394, 452, 453, 460, 481, 558, 563, 566
pmix_proc_state_t, [26](#), 26, 54, 59, 481, 521
pmix_proc_t, 20, [21](#), 21–25, 54, 67, 81, 95, 112, 121–123, 140, 150, 151, 153, 154, 157, 158, 164,
165, 174, 175, 247, 250–252, 256, 261, 264, 267, 319, 335–338, 348, 350, 356, 360–362,
364, 365, 369, 371, 373, 375, 377, 382, 385, 390, 393, 398, 399, 401, 404, 407, 409, 412,
415, 417, 418, 420, 422, 423, 461, 463–467, 481, 579, 607, 608
pmix_query_t, 55, [101](#), 101, 103, 104, 108, 112–115, 117, 118, 393, 395, 483, 587, 589, 598
pmix_rank_t, [20](#), 20, 21, 23, 24, 54, 436, 452, 459, 481, 558, 621
pmix_regattr_t, 55, 118, 344, [345](#), 345–348, 483, 592, 594, 600
pmix_scope_t, 54, 60, [75](#), 75, 481, 521
pmix_status_t, [15](#), 15, 34, 49, 50, 53, 54, 58, 59, 124, 149, 152, 154, 157, 160, 210, 339, 343, 366,
383, 386, 387, 389, 390, 398, 411, 414, 421, 480, 494, 520, 574, 619
pmix_storage_access_type_t, 55, [476](#), 476, 615
pmix_storage_accessibility_t, 55, [475](#), 475, 615
pmix_storage_medium_t, 55, [474](#), 474, 475, 615
pmix_storage_persistence_t, 55, [475](#), 475, 615
pmix_topology_t, 55, 203, [204](#), 204, 205, 208, 209, 594
pmix_value_t, 9, [29](#), 29–34, 54, 57, 58, 74, 75, 80, 81, 84, 482, 597

Index of Constants

PMIX_ALLOC_DIRECTIVE, [55](#)
PMIX_ALLOC_EXTEND, [223](#)
PMIX_ALLOC_EXTERNAL, [223](#)
PMIX_ALLOC_NEW, [223](#)
PMIX_ALLOC_REAQUIRE, [223](#)
PMIX_ALLOC_RELEASE, [223](#)
PMIX_APP, [54](#)
PMIX_APP_WILDCARD, [14](#)
PMIX_BOOL, [54](#)
PMIX_BYTE, [54](#)
PMIX_BYTE_OBJECT, [54](#)
PMIX_COMMAND, [54](#)
PMIX_COMPRESSED_BYTE_OBJECT, [55](#)
PMIX_COMPRESSED_STRING, [55](#)
PMIX_COORD, [55](#)
PMIX_COORD_LOGICAL_VIEW, [282](#)
PMIX_COORD_PHYSICAL_VIEW, [282](#)
PMIX_COORD_VIEW_UNDEF, [282](#)
PMIX_CPUBIND_PROCESS, [208](#)
PMIX_CPUBIND_THREAD, [208](#)
PMIX_DATA_ARRAY, [54](#)
PMIX_DATA_RANGE, [54](#)
PMIX_DATA_TYPE, [54](#)
PMIX_DATA_TYPE_MAX, [55](#)
PMIX_DEBUGGER_RELEASE, [458](#)
PMIX_DEVICE_DIST, [55](#)
PMIX_DEVTYPE, [55](#)
PMIX_DEVTYPE_BLOCK, [211](#)
PMIX_DEVTYPE_COPROC, [211](#)
PMIX_DEVTYPE_DMA, [211](#)
PMIX_DEVTYPE_GPU, [211](#)
PMIX_DEVTYPE_NETWORK, [211](#)
PMIX_DEVTYPE_OPENFABRICS, [211](#)
PMIX_DEVTYPE_UNKNOWN, [211](#)
PMIX_DOUBLE, [54](#)
PMIX_ENDPOINT, [55](#)
PMIX_ENVAR, [55](#)
PMIX_ERR_BAD_PARAM, [16](#)

PMIX_ERR_COMM_FAILURE, [16](#)
PMIX_ERR_CONFLICTING_CLEANUP_DIRECTIVES, [229](#)
PMIX_ERR_DUPLICATE_KEY, [133](#)
PMIX_ERR_EMPTY, [16](#)
PMIX_ERR_EVENT_REGISTRATION, [151](#)
PMIX_ERR_EXISTS, [15](#)
PMIX_ERR_EXISTS_OUTSIDE_SCOPE, [15](#)
PMIX_ERR_INIT, [16](#)
PMIX_ERR_INVALID_CRED, [15](#)
PMIX_ERR_INVALID_OPERATION, [16](#)
PMIX_ERR_IOF_COMPLETE, [448](#)
PMIX_ERR_IOF_FAILURE, [448](#)
PMIX_ERR_JOB_ABORTED, [458](#)
PMIX_ERR_JOB_ABORTED_BY_SIG, [458](#)
PMIX_ERR_JOB_ABORTED_BY_SYS_EVENT, [458](#)
PMIX_ERR_JOB_ALLOC_FAILED, [187](#)
PMIX_ERR_JOB_APP_NOT_EXECUTABLE, [187](#)
PMIX_ERR_JOB_CANCELED, [458](#)
PMIX_ERR_JOB_EXE_NOT_FOUND, [187](#)
PMIX_ERR_JOB_FAILED_TO_LAUNCH, [187](#)
PMIX_ERR_JOB_FAILED_TO_MAP, [187](#)
PMIX_ERR_JOB_INSUFFICIENT_RESOURCES, [187](#)
PMIX_ERR_JOB_KILLED_BY_CMD, [458](#)
PMIX_ERR_JOB_NO_EXE_SPECIFIED, [187](#)
PMIX_ERR_JOB_NON_ZERO_TERM, [458](#)
PMIX_ERR_JOB_SENSOR_BOUND_EXCEEDED, [458](#)
PMIX_ERR_JOB_SYS_OP_FAILED, [187](#)
PMIX_ERR_JOB_TERM_WO_SYNC, [458](#)
PMIX_ERR_JOB_WDIR_NOT_FOUND, [187](#)
PMIX_ERR_LOST_CONNECTION, [16](#)
PMIX_ERR_NO_PERMISSIONS, [16](#)
PMIX_ERR_NOMEM, [16](#)
PMIX_ERR_NOT_FOUND, [16](#)
PMIX_ERR_NOT_SUPPORTED, [16](#)
PMIX_ERR_OUT_OF_RESOURCE, [16](#)
PMIX_ERR_PACK_FAILURE, [16](#)
PMIX_ERR_PARAM_VALUE_NOT_SUPPORTED, [16](#)
PMIX_ERR_PARTIAL_SUCCESS, [17](#)
PMIX_ERR_PROC_CHECKPOINT, [229](#)
PMIX_ERR_PROC_MIGRATE, [229](#)
PMIX_ERR_PROC_RESTART, [229](#)
PMIX_ERR_PROC_TERM_WO_SYNC, [457](#)
PMIX_ERR_REPEAT_ATTR_REGISTRATION, [345](#)

PMIX_ERR_RESOURCE_BUSY, [16](#)
PMIX_ERR_TIMEOUT, [16](#)
PMIX_ERR_TYPE_MISMATCH, [15](#)
PMIX_ERR_UNKNOWN_DATA_TYPE, [15](#)
PMIX_ERR_UNPACK_FAILURE, [16](#)
PMIX_ERR_UNPACK_INADEQUATE_SPACE, [15](#)
PMIX_ERR_UNPACK_READ_PAST_END_OF_BUFFER, [16](#)
PMIX_ERR_UNREACH, [16](#)
PMIX_ERR_WOULD_BLOCK, [15](#)
PMIX_ERROR, [15](#)
PMIX_EVENT_ACTION_COMPLETE, [160](#)
PMIX_EVENT_ACTION_DEFERRED, [160](#)
PMIX_EVENT_JOB_END, [457](#)
PMIX_EVENT_JOB_START, [457](#)
PMIX_EVENT_NO_ACTION_TAKEN, [160](#)
PMIX_EVENT_NODE_DOWN, [152](#)
PMIX_EVENT_NODE_OFFLINE, [152](#)
PMIX_EVENT_PARTIAL_ACTION_TAKEN, [160](#)
PMIX_EVENT_PROC_TERMINATED, [457](#)
PMIX_EVENT_SESSION_END, [457](#)
PMIX_EVENT_SESSION_START, [457](#)
PMIX_EVENT_SYS_BASE, [152](#)
PMIX_EVENT_SYS_OTHER, [152](#)
PMIX_EXTERNAL_ERR_BASE, [17](#)
PMIX_FABRIC_REQUEST_INFO, [283](#)
PMIX_FABRIC_UPDATE_ENDPOINTS, [276](#)
PMIX_FABRIC_UPDATE_INFO, [283](#)
PMIX_FABRIC_UPDATE_PENDING, [276](#)
PMIX_FABRIC_UPDATED, [276](#)
PMIX_FLOAT, [54](#)
PMIX_FWD_ALL_CHANNELS, [448](#)
PMIX_FWD_NO_CHANNELS, [448](#)
PMIX_FWD_STDDIAG_CHANNEL, [448](#)
PMIX_FWD_STDERR_CHANNEL, [448](#)
PMIX_FWD_STDIN_CHANNEL, [448](#)
PMIX_FWD_STDOUT_CHANNEL, [448](#)
PMIX_GEOMETRY, [55](#)
PMIX_GLOBAL, [75](#)
PMIX_GROUP_ACCEPT, [270](#)
PMIX_GROUP_CONSTRUCT, [422](#)
PMIX_GROUP_CONSTRUCT_ABORT, [250](#)
PMIX_GROUP_CONSTRUCT_COMPLETE, [250](#)
PMIX_GROUP_CONTEXT_ID_ASSIGNED, [250](#)

PMIX_GROUP_DECLINE, [270](#)
PMIX_GROUP_DESTRUCT, [422](#)
PMIX_GROUP_INVITE_ACCEPTED, [250](#)
PMIX_GROUP_INVITE_DECLINED, [250](#)
PMIX_GROUP_INVITE_FAILED, [250](#)
PMIX_GROUP_INVITED, [249](#)
PMIX_GROUP_LEADER_FAILED, [250](#)
PMIX_GROUP_LEADER_SELECTED, [250](#)
PMIX_GROUP_LEFT, [249](#)
PMIX_GROUP_MEMBER_FAILED, [249](#)
PMIX_GROUP_MEMBERSHIP_UPDATE, [250](#)
PMIX_INFO, [54](#)
PMIX_INFO_ARRAY_END, [41](#)
PMIX_INFO_DIR_RESERVED, [41](#)
PMIX_INFO_DIRECTIVES, [54](#)
PMIX_INFO_REQD, [41](#)
PMIX_INFO_REQD_PROCESSED, [41](#)
PMIX_INT, [54](#)
PMIX_INT16, [54](#)
PMIX_INT32, [54](#)
PMIX_INT64, [54](#)
PMIX_INT8, [54](#)
PMIX_INTERNAL, [75](#)
PMIX_IOF_CHANNEL, [55](#)
PMIX_JCTRL_CHECKPOINT, [229](#)
PMIX_JCTRL_CHECKPOINT_COMPLETE, [229](#)
PMIX_JCTRL_PREEMPT_ALERT, [229](#)
PMIX_JOB_STATE, [55](#)
PMIX_JOB_STATE_AWAITING_ALLOC, [29](#)
PMIX_JOB_STATE_CONNECTED, [29](#)
PMIX_JOB_STATE_LAUNCH_UNDERWAY, [29](#)
PMIX_JOB_STATE_RUNNING, [29](#)
PMIX_JOB_STATE_SUSPENDED, [29](#)
PMIX_JOB_STATE_TERMINATED, [29](#)
PMIX_JOB_STATE_TERMINATED_WITH_ERROR, [29](#)
PMIX_JOB_STATE_UNDEF, [29](#)
PMIX_JOB_STATE_UNTERMINATED, [29](#)
PMIX_KVAL, [54](#)
PMIX_LAUNCH_COMPLETE, [457](#)
PMIX_LAUNCHER_READY, [443](#)
PMIX_LINK_DOWN, [283](#)
PMIX_LINK_STATE, [55](#)
PMIX_LINK_STATE_UNKNOWN, [283](#)

PMIX_LINK_UP, [283](#)
PMIX_LOCAL, [75](#)
PMIX_LOCALITY_NONLOCAL, [206](#)
PMIX_LOCALITY_SHARE_CORE, [206](#)
PMIX_LOCALITY_SHARE_HWTHREAD, [206](#)
PMIX_LOCALITY_SHARE_L1CACHE, [206](#)
PMIX_LOCALITY_SHARE_L2CACHE, [206](#)
PMIX_LOCALITY_SHARE_L3CACHE, [206](#)
PMIX_LOCALITY_SHARE_NODE, [206](#)
PMIX_LOCALITY_SHARE_NUMA, [206](#)
PMIX_LOCALITY_SHARE_PACKAGE, [206](#)
PMIX_LOCALITY_UNKNOWN, [206](#)
PMIX_LOCTYPE, [55](#)
PMIX_MAX_KEYLEN, [14](#)
PMIX_MAX_NSLEN, [14](#)
PMIX_MODEL_DECLARED, [67](#)
PMIX_MODEL_RESOURCES, [67](#)
PMIX_MONITOR_FILE_ALERT, [236](#)
PMIX_MONITOR_HEARTBEAT_ALERT, [236](#)
PMIX_OPENMP_PARALLEL_ENTERED, [67](#)
PMIX_OPENMP_PARALLEL_EXITED, [67](#)
PMIX_OPERATION_IN_PROGRESS, [16](#)
PMIX_OPERATION_SUCCEEDED, [16](#)
PMIX_PDATA, [54](#)
PMIX_PERSIST, [54](#)
PMIX_PERSIST_APP, [134](#)
PMIX_PERSIST_FIRST_READ, [134](#)
PMIX_PERSIST_INDEF, [134](#)
PMIX_PERSIST_INVALID, [134](#)
PMIX_PERSIST_PROC, [134](#)
PMIX_PERSIST_SESSION, [134](#)
PMIX_PID, [54](#)
PMIX_POINTER, [54](#)
PMIX_PROC, [54](#)
PMIX_PROC_CPUSET, [55](#)
PMIX_PROC_INFO, [54](#)
PMIX_PROC_NAMESPACE, [54](#)
PMIX_PROC_RANK, [54](#)
PMIX_PROC_STATE, [54](#)
PMIX_PROC_STATE_ABORTED, [26](#)
PMIX_PROC_STATE_ABORTED_BY_SIG, [26](#)
PMIX_PROC_STATE_CALLED_ABORT, [26](#)
PMIX_PROC_STATE_CANNOT_RESTART, [26](#)

[PMIX_PROC_STATE_COMM_FAILED, 26](#)
[PMIX_PROC_STATE_CONNECTED, 26](#)
[PMIX_PROC_STATE_ERROR, 26](#)
[PMIX_PROC_STATE_FAILED_TO_LAUNCH, 26](#)
[PMIX_PROC_STATE_FAILED_TO_START, 26](#)
[PMIX_PROC_STATE_HEARTBEAT_FAILED, 26](#)
[PMIX_PROC_STATE_KILLED_BY_CMD, 26](#)
[PMIX_PROC_STATE_LAUNCH_UNDERWAY, 26](#)
[PMIX_PROC_STATE_MIGRATING, 26](#)
[PMIX_PROC_STATE_PREPPED, 26](#)
[PMIX_PROC_STATE_RESTART, 26](#)
[PMIX_PROC_STATE_RUNNING, 26](#)
[PMIX_PROC_STATE_SENSOR_BOUND_EXCEEDED, 26](#)
[PMIX_PROC_STATE_TERM_NON_ZERO, 26](#)
[PMIX_PROC_STATE_TERM_WO_SYNC, 26](#)
[PMIX_PROC_STATE_TERMINATE, 26](#)
[PMIX_PROC_STATE_TERMINATED, 26](#)
[PMIX_PROC_STATE_UNDEF, 26](#)
[PMIX_PROC_STATE_UNTERMINATED, 26](#)
[PMIX_PROCESS_SET_DEFINE, 246](#)
[PMIX_PROCESS_SET_DELETE, 246](#)
[PMIX_QUERY, 55](#)
[PMIX_RANGE_CUSTOM, 134](#)
[PMIX_RANGE_GLOBAL, 134](#)
[PMIX_RANGE_INVALID, 134](#)
[PMIX_RANGE_LOCAL, 134](#)
[PMIX_RANGE_NAMESPACE, 134](#)
[PMIX_RANGE_PROC_LOCAL, 134](#)
[PMIX_RANGE_RM, 134](#)
[PMIX_RANGE_SESSION, 134](#)
[PMIX_RANGE_UNDEF, 134](#)
[PMIX_RANK_INVALID, 21](#)
[PMIX_RANK_LOCAL_NODE, 20](#)
[PMIX_RANK_LOCAL_PEERS, 21](#)
[PMIX_RANK_UNDEF, 20](#)
[PMIX_RANK_VALID, 21](#)
[PMIX_RANK_WILDCARD, 20](#)
[PMIX_READY_FOR_DEBUG, 458](#)
[PMIX_REGATTR, 55](#)
[PMIX_REGEX, 55](#)
[PMIX_REMOTE, 75](#)
[PMIX_SCOPE, 54](#)
[PMIX_SCOPE_UNDEF, 75](#)

[PMIX_SIZE, 54](#)
[PMIX_STATUS, 54](#)
[PMIX_STOR_ACCESS, 55](#)
[PMIX_STOR_ACCESS_TYPE, 55](#)
[PMIX_STOR_MEDIUM, 55](#)
[PMIX_STOR_PERSIST, 55](#)
[PMIX_STORAGE_ACCESS_RD, 476](#)
[PMIX_STORAGE_ACCESS_RDWR, 476](#)
[PMIX_STORAGE_ACCESS_WR, 476](#)
[PMIX_STORAGE_ACCESSIBILITY_CLUSTER, 475](#)
[PMIX_STORAGE_ACCESSIBILITY_JOB, 475](#)
[PMIX_STORAGE_ACCESSIBILITY_NODE, 475](#)
[PMIX_STORAGE_ACCESSIBILITY_RACK, 475](#)
[PMIX_STORAGE_ACCESSIBILITY_REMOTE, 475](#)
[PMIX_STORAGE_ACCESSIBILITY_SESSION, 475](#)
[PMIX_STORAGE_MEDIUM_HDD, 474](#)
[PMIX_STORAGE_MEDIUM_NVME, 474](#)
[PMIX_STORAGE_MEDIUM_PMEM, 474](#)
[PMIX_STORAGE_MEDIUM_RAM, 474](#)
[PMIX_STORAGE_MEDIUM_SSD, 474](#)
[PMIX_STORAGE_MEDIUM_TAPE, 474](#)
[PMIX_STORAGE_MEDIUM_UNKNOWN, 474](#)
[PMIX_STORAGE_PERSISTENCE_ARCHIVE, 476](#)
[PMIX_STORAGE_PERSISTENCE_JOB, 475](#)
[PMIX_STORAGE_PERSISTENCE_NODE, 475](#)
[PMIX_STORAGE_PERSISTENCE_PROJECT, 476](#)
[PMIX_STORAGE_PERSISTENCE_SCRATCH, 475](#)
[PMIX_STORAGE_PERSISTENCE_SESSION, 475](#)
[PMIX_STORAGE_PERSISTENCE_TEMPORARY, 475](#)
[PMIX_STRING, 54](#)
[PMIX_SUCCESS, 15](#)
[PMIX_TIME, 54](#)
[PMIX_TIMEVAL, 54](#)
[PMIX_TOPO, 55](#)
[PMIX_UINT, 54](#)
[PMIX_UINT16, 54](#)
[PMIX_UINT32, 54](#)
[PMIX_UINT64, 54](#)
[PMIX_UINT8, 54](#)
[PMIX_UNDEF, 54](#)
[PMIX_VALUE, 54](#)

PMIX_BUFFER
Removed, 614

PMIX_CONNECT_REQUESTED
Deprecated, [611](#)

PMIX_DEBUG_WAITING_FOR_NOTIFY
Deprecated, [617](#)

PMIX_ERR_DATA_VALUE_NOT_FOUND
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_DEBUGGER_RELEASE
Deprecated, [611](#)

PMIX_ERR_HANDSHAKE_FAILED
Deprecated, [590](#)
Removed, [611](#)

PMIX_ERR_IN_ERRNO
Deprecated, [590](#)
Removed, [611](#)

PMIX_ERR_INVALID_ARG
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_ARGS
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_KEY
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_KEY_LENGTH
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_KEYVALP
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_LENGTH
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_NAMESPACE
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_NUM_ARGS
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_NUM_PARSED
Deprecated, [590](#)
Removed, [612](#)

PMIX_ERR_INVALID_SIZE

Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_INVALID_TERMINATION
Deprecated, [611](#)
PMIX_ERR_INVALID_VAL
Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_INVALID_VAL_LENGTH
Deprecated, [590](#)
Removed, [611](#)
PMIX_ERR_JOB_TERMINATED
Deprecated, [611](#)
PMIX_ERR_LOST_CONNECTION_TO_CLIENT
Deprecated, [611](#)
PMIX_ERR_LOST_CONNECTION_TO_SERVER
Deprecated, [611](#)
PMIX_ERR_LOST_PEER_CONNECTION
Deprecated, [611](#)
PMIX_ERR_NODE_DOWN
Deprecated, [611](#)
PMIX_ERR_NODE_OFFLINE
Deprecated, [611](#)
PMIX_ERR_NOT_IMPLEMENTED
Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_PACK_MISMATCH
Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_PROC_ABORTED
Deprecated, [611](#)
PMIX_ERR_PROC_ABORTING
Deprecated, [611](#)
PMIX_ERR_PROC_ENTRY_NOT_FOUND
Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_PROC_REQUESTED_ABORT
Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_READY_FOR_HANDSHAKE
Deprecated, [590](#)
Removed, [611](#)
PMIX_ERR_SERVER_FAILED_REQUEST
Deprecated, [590](#)

Removed, [612](#)
PMIX_ERR_SERVER_NOT_AVAIL
Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_SILENT
Deprecated, [590](#)
Removed, [612](#)
PMIX_ERR_SYS_OTHER
Deprecated, [611](#)
PMIX_EXISTS
Deprecated, [611](#)
PMIX_GDS_ACTION_COMPLETE
Deprecated, [590](#)
Removed, [612](#)
PMIX_INFO_ARRAY
Deprecated, [585](#)
PMIX_MODEX
Deprecated, [585](#)
PMIX_NOTIFY_ALLOC_COMPLETE
Deprecated, [590](#)
Removed, [612](#)
PMIX_PROC_HAS_CONNECTED
Deprecated, [611](#)
PMIX_PROC_TERMINATED
Deprecated, [611](#)

Unofficial Draft

Index of Environmental Variables

PMIX_KEEPALIVE_PIPE, [431](#), [441](#), [442](#), [610](#)

PMIX_LAUNCHER_RNDZ_FILE, [427](#), [431](#), [610](#)

PMIX_LAUNCHER_RNDZ_URI, [431](#), [441](#), [442](#), [563](#), [610](#)

Unofficial Draft

Index of Attributes

PMIX_ACCESS_GRPIDS, 127, 129, 131, 132, [133](#), 601
PMIX_ACCESS_PERMISSIONS, 127, 129, 131, 132, [133](#), 601
PMIX_ACCESS_USERIDS, 127, 129, 130, 132, [133](#), 601
PMIX_ADD_ENVAR, 179, 185, [191](#), 569
PMIX_ADD_HOST, 177, 183, [188](#), 379
PMIX_ADD_HOSTFILE, 177, 183, [188](#), 379
PMIX_ALL_CLONES_PARTICIPATE, 122, 124, [125](#), 196, 199, 200, 202, 597
PMIX_ALLOC_BANDWIDTH, 180, 186, 217, 220, [222](#), 222, 341, 403
PMIX_ALLOC_CPU_LIST, 179, 185, 217, 220, [222](#), 402
PMIX_ALLOC_FABRIC, [217](#), 220, [222](#), 340, 402, 612
PMIX_ALLOC_FABRIC_ENDPTS, 180, 186, 217, 218, 220, 221, [222](#), 222, 340, 341, 402, 613
PMIX_ALLOC_FABRIC_ENDPTS_NODE, 180, 186, 218, 221, [222](#), 341, 613
PMIX_ALLOC_FABRIC_ID, [217](#), 220, [222](#), 222, 340, 402, 612
PMIX_ALLOC_FABRIC_PLANE, 180, 186, 217, 220, 221, [222](#), 222, 341, 403, 612
PMIX_ALLOC_FABRIC_QOS, 180, 186, 217, 220, 221, [222](#), 222, 341, 403, 612
PMIX_ALLOC_FABRIC_SEC_KEY, [217](#), 218, 220, 221, [222](#), 222, 341, 403, 613
PMIX_ALLOC_FABRIC_TYPE, 180, 186, 217, 220, 221, [222](#), 222, 340, 341, 402, 403, 612
PMIX_ALLOC_ID, [218](#), [221](#), 402, 589, 590, 601
PMIX_ALLOC_MEM_SIZE, 179, 186, 217, 220, [222](#), 402
PMIX_ALLOC_NODE_LIST, 179, 185, 217, 220, [221](#), 402
PMIX_ALLOC_NUM_CPU_LIST, 179, 185, 217, 220, [222](#), 402
PMIX_ALLOC_NUM_CPUS, 179, 185, 216, 220, [221](#), 402
PMIX_ALLOC_NUM_NODES, 179, 185, 216, 219, [221](#), 402
PMIX_ALLOC_QUEUE, [104](#), [109](#), [111](#), 179, 185, [221](#), 394, 601
PMIX_ALLOC_REQ_ID, [216](#), [219](#), [221](#), 589
PMIX_ALLOC_TIME, 179, 185, 216, 220, [222](#), 402
PMIX_ALLOCATED_NODELIST, [88](#), 314
PMIX_ANL_MAP, [89](#), [90](#), 316
PMIX_APP_ARGV, [90](#), [91](#), 317, 601
PMIX_APP_INFO, [80](#), [83](#), [86](#), [91](#), [102](#), [107](#), 317, 318, 551
PMIX_APP_INFO_ARRAY, [313](#), [316](#), [322](#), 323, 327, 600
PMIX_APP_MAP_REGEX, [91](#), 317
PMIX_APP_MAP_TYPE, [91](#), 317
PMIX_APP_RANK, [93](#), 319
PMIX_APP_SIZE, [91](#), 317, 327
PMIX_APPEND_ENVAR, 179, 185, [191](#), 569
PMIX_APPLDR, [91](#), 317, 327
PMIX_APPNUM, [80](#), [83](#), [86](#), [87](#), [91](#), [92](#), [102](#), [107](#), [313](#), [316](#), [318](#), [319](#), [322](#), [327](#), 551, 600

PMIX_ATTR_UNDEF, [5](#)
 PMIX_AVAIL_PHYS_MEMORY, [95](#), 112, 319
 PMIX_BINDTO, [177](#), 183, [188](#), 316, 379
 PMIX_BREAKPOINT, [436](#), 452, 458, [459](#), 459, 558, 621
 PMIX_CLEANUP_EMPTY, [225](#), 228, [231](#)
 PMIX_CLEANUP_IGNORE, [225](#), 228, [231](#)
 PMIX_CLEANUP_LEAVE_TOPDIR, [225](#), 228, [231](#)
 PMIX_CLEANUP_RECURSIVE, [225](#), 228, [231](#)
 PMIX_CLIENT_ATTRIBUTES, [103](#), 108, [113](#), 118, 452, 592, 599
 PMIX_CLIENT_AVG_MEMORY, [105](#), 110, [112](#)
 PMIX_CLIENT_FUNCTIONS, [103](#), 104, 109, 111, [113](#), 118, 598, 599
 PMIX_CLUSTER_ID, [87](#), 314
 PMIX_CMD_LINE, [90](#), 601, 617
 PMIX_COLLECT_DATA, [73](#), 120, 121, 123, [124](#), 366
 PMIX_COLLECT_GENERATED_JOB_INFO, [121](#), 123, 124, [125](#), 274, 366, 597
 PMIX_COLLECTIVE_ALGO, 586
 PMIX_CONNECT_MAX_RETRIES, [431](#), 462
 PMIX_CONNECT_RETRY_DELAY, [432](#), 462
 PMIX_CONNECT_SYSTEM_FIRST, [429](#), [431](#), 462, 465
 PMIX_CONNECT_TO_SYSTEM, [429](#), [431](#), 462, 465
 PMIX_COSPAWN_APP, [180](#), 186, [460](#), 557, 561
 PMIX_CPU_LIST, [178](#), 184, [189](#), 381
 PMIX_CPUS_PER_PROC, [178](#), 184, [189](#), 381
 PMIX_CPuset, [94](#), 207, 321, 354
 PMIX_CPuset_BITMAP, [94](#), 321, 602
 PMIX_CRED_TYPE, [299](#), 410
 PMIX_CREDENTIAL, [94](#), 396
 PMIX_CRYPT_KEY, [299](#), 316
 PMIX_DAEMON_MEMORY, [105](#), 110, [112](#)
 PMIX_DATA_SCOPE, [79](#), 83, [84](#)
 PMIX_DEBUG_DAEMONS_PER_NODE, [380](#), 454, [460](#), 460, 556, 558, 562, 563, 566, 604
 PMIX_DEBUG_DAEMONS_PER_PROC, [380](#), 454, [460](#), 460, 556, 558, 562, 563, 566, 603
 PMIX_DEBUG_STOP_IN_APP, [436](#), 451, 452, [459](#), 558, 617, 621
 PMIX_DEBUG_STOP_IN_INIT, [436](#), 440, 442, 451, 452, 455, [459](#), 556, 562
 PMIX_DEBUG_STOP_ON_EXEC, [436](#), 442, 451, 452, [459](#), 556, 562
 PMIX_DEBUG_TARGET, [380](#), 452, 454, 455, [459](#), 459, 460, 556–558, 562, 563, 566, 603, 604,
 [612](#)
 PMIX_DEBUGGER_DAEMONS, [380](#), 453–455, [459](#), 553, 557, 563, 566
 PMIX_DEVICE_DISTANCES, [214](#), 290, 321, 607
 PMIX_DEVICE_ID, [214](#), 274, 275, 285, 288, 290, 334, 605–607
 PMIX_DEVICE_TYPE, [214](#), 607
 PMIX_DISPLAY_MAP, [177](#), 183, [188](#), 379
 PMIX_EMBED_BARRIER, [69](#), 69

PMIX_ENUM_VALUE, [346](#), [346](#), [592](#), [600](#)
PMIX_ENVARS_HARVESTED, [180](#), [186](#), [191](#), [619](#)
PMIX_EVENT_ACTION_TIMEOUT, [153](#), [158](#)
PMIX_EVENT_AFFECTED_PROC, [151](#), [153](#), [158](#), [457](#), [557](#)
PMIX_EVENT_AFFECTED_PROCS, [151](#), [153](#), [158](#), [457](#)
PMIX_EVENT_BASE, [66](#), [307](#), [463](#)
PMIX_EVENT_CUSTOM_RANGE, [150](#), [153](#), [158](#)
PMIX_EVENT_DO_NOT_CACHE, [153](#), [158](#)
PMIX_EVENT_HDLR_AFTER, [150](#), [152](#), [577](#)
PMIX_EVENT_HDLR_APPEND, [150](#), [153](#), [577](#)
PMIX_EVENT_HDLR_BEFORE, [150](#), [152](#), [577](#)
PMIX_EVENT_HDLR_FIRST, [150](#), [152](#), [577](#)
PMIX_EVENT_HDLR_FIRST_IN_CATEGORY, [150](#), [152](#), [577](#)
PMIX_EVENT_HDLR_LAST, [150](#), [152](#), [577](#)
PMIX_EVENT_HDLR_LAST_IN_CATEGORY, [150](#), [152](#), [577](#)
PMIX_EVENT_HDLR_NAME, [150](#), [152](#), [577](#)
PMIX_EVENT_HDLR_PREPEND, [150](#), [152](#)
PMIX_EVENT_NON_DEFAULT, [153](#), [158](#)
PMIX_EVENT_PROXY, [153](#), [158](#)
PMIX_EVENT_RETURN_OBJECT, [150](#), [153](#)
PMIX_EVENT_SILENT_TERMINATION, [180](#), [186](#), [191](#)
PMIX_EVENT_TERMINATE_JOB, [153](#), [158](#)
PMIX_EVENT_TERMINATE_NODE, [153](#), [158](#)
PMIX_EVENT_TERMINATE_PROC, [153](#), [158](#)
PMIX_EVENT_TERMINATE_SESSION, [153](#), [158](#)
PMIX_EVENT_TEXT_MESSAGE, [153](#), [158](#)
PMIX_EVENT_TIMESTAMP, [153](#), [190](#), [191](#), [435–437](#), [456](#), [457](#), [557](#), [563](#), [567](#), [609](#), [610](#)
PMIX_EXEC_AGENT, [440](#), [443](#), [456](#), [603](#)
PMIX_EXIT_CODE, [93](#), [190](#), [191](#), [435–437](#), [456](#), [557](#), [563](#), [567](#), [610](#)
PMIX_EXTERNAL_PROGRESS, [65](#), [307](#), [310](#), [463](#), [602](#)
PMIX_FABRIC_COORDINATES, [288](#), [605](#)
PMIX_FABRIC_COST_MATRIX, [284](#), [287](#), [604](#)
PMIX_FABRIC_DEVICE, [274](#), [285](#), [288](#), [288](#), [606](#)
PMIX_FABRIC_DEVICE_ADDRESS, [275](#), [285](#), [289](#), [606](#)
PMIX_FABRIC_DEVICE_BUS_TYPE, [275](#), [286](#), [289](#), [606](#)
PMIX_FABRIC_DEVICE_COORDINATES, [275](#), [289](#), [606](#)
PMIX_FABRIC_DEVICE_DRIVER, [275](#), [285](#), [289](#), [606](#)
PMIX_FABRIC_DEVICE_FIRMWARE, [275](#), [285](#), [289](#), [606](#)
PMIX_FABRIC_DEVICE_INDEX, [275](#), [288](#), [424](#), [606](#)
PMIX_FABRIC_DEVICE_MTU, [276](#), [285](#), [289](#), [606](#)
PMIX_FABRIC_DEVICE_NAME, [274](#), [275](#), [285](#), [288](#), [334](#), [606](#)
PMIX_FABRIC_DEVICE_PCI_DEVID, [275](#), [286](#), [289](#), [289](#), [607](#)
PMIX_FABRIC_DEVICE_SPEED, [276](#), [286](#), [289](#), [606](#)

PMIX_FABRIC_DEVICE_STATE, 276, 286, **289**, 607
PMIX_FABRIC_DEVICE_TYPE, 276, 286, **289**, 607
PMIX_FABRIC_DEVICE_VENDOR, 275, 285, **289**, 606
PMIX_FABRIC_DEVICE_VENDORID, 275, **289**, 606
PMIX_FABRIC_DEVICES, 274, **288**
PMIX_FABRIC_DIMS, 284, **288**, 605
PMIX_FABRIC_ENDPT, **289**, 605
PMIX_FABRIC_GROUPS, 284, **287**, 604
PMIX_FABRIC_IDENTIFIER, 284, **287**, 291, 424, 604
PMIX_FABRIC_INDEX, 283, 287, **288**, 604
PMIX_FABRIC_NUM_DEVICES, 284, **288**, 605
PMIX_FABRIC_PLANE, 284, 285, **287**, 288, 291, 292, 424, 605
PMIX_FABRIC_SHAPE, 285, **288**, 605
PMIX_FABRIC_SHAPE_STRING, 285, **288**, 605
PMIX_FABRIC_SWITCH, **287**, 290, 605, 606
PMIX_FABRIC_VENDOR, 284, **287**, 291, 423, 604
PMIX_FIRST_ENVAR, 179, 185, **191**, 610
PMIX_FORKEXEC_AGENT, 440, 442, **443**, 455, 603
PMIX_FWD_STDDIAG, 434, 439, **443**, 589
PMIX_FWD_STDERR, 380, 397, 433, 439, **443**, 445, 557, 561, 567
PMIX_FWD_STDIN, 379, 397, 433, 439, **442**, 445
PMIX_FWD_STDOUT, 380, 397, 433, 439, 442, **443**, 445, 557, 561, 567
PMIX_GET_POINTER_VALUES, 80, 81, 83, **84**, 597
PMIX_GET_REFRESH_CACHE, 77, 80, 83, **84**, 597
PMIX_GET_STATIC_VALUES, 79–81, **84**, 84, 597
PMIX_GLOBAL_RANK, 89, **93**, 316, 319, 552
PMIX_GROUP_ASSIGN_CONTEXT_ID, **251**, 253, 257, 262, 265, 421, 422, 608
PMIX_GROUP_CONTEXT_ID, **251**, 422, 609
PMIX_GROUP_ENDPT_DATA, **251**, 421, 422, 609
PMIX_GROUP_FT_COLLECTIVE, **251**, 253, 257, 262, 265, 608
PMIX_GROUP_ID, **250**, 250, 422, 579, 608
PMIX_GROUP_LEADER, **251**, 253, 254, 257, 263, 268, 608
PMIX_GROUP_LOCAL_ONLY, **251**, 253, 257, 421, 608
PMIX_GROUP_MEMBERSHIP, **251**, 254, 422
PMIX_GROUP_NAMES, **251**, 609
PMIX_GROUP_NOTIFY_TERMINATION, **251**, 253, 254, 257, 259, 262, 265, 608
PMIX_GROUP_OPTIONAL, **251**, 253, 254, 257, 262, 265, 421, 608
PMIX_GRPID, 127, 129, 130, 132, 136, 138, 143, 144, 216, 219, 224, 227, 232, 234, 238, 241,
297, 299, 301, 302, 371, 372, 374, 376, 378, 388, 394, 396, **398**, 399, 402, 405, 408, 410,
413, 414, 416, 419
PMIX_HOMOGENEOUS_SYSTEM, 307, **310**, 600
PMIX_HOST, 176, 183, **188**, 378
PMIX_HOST_ATTRIBUTES, 103, 108, **113**, 118, 119, 452, 592, 599

PMIX_HOST_FUNCTIONS, 103, 104, 109, 111, **113**, 118, 598, 599
PMIX_HOSTFILE, 177, 183, **188**, 379
PMIX_HOSTNAME, 80, 83, 87, 90, **94**, 94, 95, 102, 104, 105, 107, 109–112, 274, 275, 285, 286,
289, 314, 318, 320, 321, 323, 334, 394, 453, 460, 551, 552, 558, 563, 600, 607
PMIX_HOSTNAME_ALIASES, **94**, 318, 602
PMIX_HOSTNAME_KEEP_FQDN, **88**, 316, 602
PMIX_IMMEDIATE, 77, 79, 82, **84**, 95
PMIX_INDEX_ARGV, 178, 184, **189**, 380
PMIX_IOF_BUFFERING_SIZE, 416, 434, 439, **449**, 469, 472
PMIX_IOF_BUFFERING_TIME, 416, 434, 439, **449**, 469, 472
PMIX_IOF_CACHE_SIZE, 416, 434, 439, **449**, 469, 472
PMIX_IOF_COMPLETE, 417, 447, **449**, 450, 473, 486, 603
PMIX_IOF_COPY, 445, **450**, 603
PMIX_IOF_DROP_NEWEST, 416, 434, 439, **449**, 469, 472
PMIX_IOF_DROP_OLDEST, 416, 434, 439, **449**, 469, 472
PMIX_IOF_FILE_ONLY, 435, 437, 446, **450**, 620
PMIX_IOF_FILE_PATTERN, 435, 446, **450**, 620
PMIX_IOF_LOCAL_OUTPUT, 307, **449**, 463, 620
PMIX_IOF_MERGE_STDERR_STDOUT, 435, 446, **449**, 450, 620
PMIX_IOF_OUTPUT_RAW, 434, **449**, 620
PMIX_IOF_OUTPUT_TO_DIRECTORY, 435, 437, 446, **450**, 620
PMIX_IOF_OUTPUT_TO_FILE, 435, 437, 446, **450**, 620
PMIX_IOF_PUSH_STDIN, 447, **449**, 472, 603
PMIX_IOF_RANK_OUTPUT, 434, 446, **449**, 620
PMIX_IOF_REDIRECT, 445, **450**, 603
PMIX_IOF_TAG_OUTPUT, 434, 440, 446, **449**, 469
PMIX_IOF_TIMESTAMP_OUTPUT, 434, 440, 446, **449**, 469
PMIX_IOF_XML_OUTPUT, 434, 440, 446, **449**, 470
PMIX_JOB_CONTINUOUS, 178, 184, **190**, 381
PMIX_JOB_CTRL_CANCEL, 225, 228, **230**, 405
PMIX_JOB_CTRL_CHECKPOINT, 225, 228, **230**, 405
PMIX_JOB_CTRL_CHECKPOINT_EVENT, 225, 228, **230**, 406
PMIX_JOB_CTRL_CHECKPOINT_METHOD, 226, 229, **230**, 406
PMIX_JOB_CTRL_CHECKPOINT_SIGNAL, 225, 228, **230**, 406
PMIX_JOB_CTRL_CHECKPOINT_TIMEOUT, 226, 228, **230**, 406
PMIX_JOB_CTRL_ID, 224, 225, 227, 228, **230**, 230, 405
PMIX_JOB_CTRL_KILL, 225, 227, **230**, 405
PMIX_JOB_CTRL_PAUSE, 224, 227, **230**, 405
PMIX_JOB_CTRL_PREEMPTIBLE, 226, 229, **230**, 406
PMIX_JOB_CTRL_PROVISION, 226, 229, **230**, 406
PMIX_JOB_CTRL_PROVISION_IMAGE, 226, 229, **230**, 406
PMIX_JOB_CTRL_RESTART, 225, 228, **230**, 405
PMIX_JOB_CTRL_RESUME, 224, 227, **230**, 405

PMIX_JOB_CTRL_SIGNAL, 225, 228, [230](#), 405
PMIX_JOB_CTRL_TERMINATE, 225, 228, [230](#), 405
PMIX_JOB_INFO, [80](#), [83](#), [86](#), 89, 102, 107
PMIX_JOB_INFO_ARRAY, [313](#), [315](#), [322](#), 323, 327, 587, 599
PMIX_JOB_NUM_APPS, [90](#), 316, 326
PMIX_JOB_RECOVERABLE, 178, 184, [190](#), 381
PMIX_JOB_SIZE, [90](#), 315, 326, 327, 551, 586, 589
PMIX_JOB_TERM_STATUS, 190, 191, 435–437, 456, [458](#), 458, 557, 563, 567, 610
PMIX_JOB_TIMEOUT, 180, 186, [191](#), 191, 381, 619, 620
PMIX_JOBID, [89](#), [313](#), 315, 322, 326, 457, 600
PMIX_LAUNCH_DIRECTIVES, 442, [443](#), 561, 604
PMIX_LAUNCHER, 426, [431](#), 432
PMIX_LAUNCHER_DAEMON, 440, [443](#), 603
PMIX_LAUNCHER_RENDEZVOUS_FILE, 427, [432](#), 602
PMIX_LOCAL_COLLECTIVE_STATUS, [124](#), 366, 383, 385, 421, 619
PMIX_LOCAL_CPUSSETS, [91](#), 319, 330
PMIX_LOCAL_PEERS, [90](#), 91, 318, 319, 328, 551
PMIX_LOCAL_PROCS, [95](#), 319
PMIX_LOCAL_RANK, [93](#), 320, 454, 455, 460, 552, 556, 557, 562, 566, 604
PMIX_LOCAL_SIZE, [91](#), 318, 551
PMIX_LOCALITY_STRING, 204, [206](#), 320, 354, 552
PMIX_LOCALLDR, [90](#), 318
PMIX_LOG_COMPLETION, [191](#), 436, 457, 610
PMIX_LOG_EMAIL, 239, 242, [244](#), 400
PMIX_LOG_EMAIL_ADDR, 239, 242, [244](#), 400
PMIX_LOG_EMAIL_MSG, 239, 242, [244](#), 400
PMIX_LOG_EMAIL_SENDER_ADDR, 239, 242, [244](#)
PMIX_LOG_EMAIL_SERVER, 239, 242, [244](#)
PMIX_LOG_EMAIL_SRVR_PORT, 239, 242, [244](#)
PMIX_LOG_EMAIL_SUBJECT, 239, 242, [244](#), 400
PMIX_LOG_GENERATE_TIMESTAMP, 238, 241, [243](#)
PMIX_LOG_GLOBAL_DATASTORE, 239, 242, [244](#)
PMIX_LOG_GLOBAL_SYSLOG, 238, 241, [243](#)
PMIX_LOG_JOB_EVENTS, [190](#), 436, 456, 610
PMIX_LOG_JOB_RECORD, 239, 242, [244](#)
PMIX_LOG_LOCAL_SYSLOG, 238, 241, [243](#)
PMIX_LOG_MSG, [243](#), 400
PMIX_LOG_ONCE, 238, 241, [243](#)
PMIX_LOG_PROC_ABNORMAL_TERMINATION, [190](#), 610
PMIX_LOG_PROC_TERMINATION, [190](#), 609
PMIX_LOG_SOURCE, 238, 241, [243](#)
PMIX_LOG_STDERR, 238, 241, [243](#), 399
PMIX_LOG_STDOUT, 238, 241, [243](#), 399

PMIX_LOG_SYSLOG, 238, 241, [243](#), 399
PMIX_LOG_SYSLOG_PRI, 238, 241, [243](#)
PMIX_LOG_TAG_OUTPUT, 239, 242, [243](#)
PMIX_LOG_TIMESTAMP, 238, 241, [243](#)
PMIX_LOG_TIMESTAMP_OUTPUT, 239, 242, [243](#)
PMIX_LOG_XML_OUTPUT, 239, 242, [243](#)
PMIX_MAPBY, 177, 183, [188](#), 188, 316, 379, 557, 562, 566
PMIX_MAX_PROCS, 87, [88](#), 88–90, 92, 314, 315, 317, 319, 345, 551, 589
PMIX_MAX_RESTARTS, 178, 185, [190](#), 381
PMIX_MAX_VALUE, [346](#), 346, 592, 600
PMIX_MERGE_STDERR_STDOUT, 178, 184, [189](#), 380
PMIX_MIN_VALUE, [346](#), 346, 592, 600
PMIX_MODEL_AFFINITY_POLICY, 66, [68](#), 572
PMIX_MODEL_CPU_TYPE, 66, [68](#), 571
PMIX_MODEL_LIBRARY_NAME, 66, [68](#), 317, 342, 571
PMIX_MODEL_LIBRARY_VERSION, 66, [68](#), 318, 342, 571
PMIX_MODEL_NUM_CPUS, 66, [68](#), 571
PMIX_MODEL_NUM_THREADS, 66, [68](#), 571
PMIX_MODEL_PHASE_NAME, [68](#), 153, 571, 572
PMIX_MODEL_PHASE_TYPE, [68](#), 153, 571
PMIX_MONITOR_APP_CONTROL, 232, 235, [236](#), 408
PMIX_MONITOR_CANCEL, 232, 235, [236](#), 408
PMIX_MONITOR_FILE, 232, 233, 235, [236](#), 408
PMIX_MONITOR_FILE_ACCESS, 232, 235, [237](#), 408
PMIX_MONITOR_FILE_CHECK_TIME, 233, 235, [237](#), 409
PMIX_MONITOR_FILE_DROPS, 233, 235, [237](#), 409
PMIX_MONITOR_FILE_MODIFY, 233, 235, [237](#), 408
PMIX_MONITOR_FILE_SIZE, 232, 235, [236](#), 408
PMIX_MONITOR_HEARTBEAT, 232, 235, [236](#), 408
PMIX_MONITOR_HEARTBEAT_DROPS, 232, 235, [236](#), 408
PMIX_MONITOR_HEARTBEAT_TIME, 232, 235, [236](#), 408
PMIX_MONITOR_ID, 232, 234, [236](#), 408
PMIX_NO_OVERSUBSCRIBE, 178, 184, [189](#), 381
PMIX_NO_PROCS_ON_HEAD, 178, 184, [189](#), 381
PMIX_NODE_INFO, 80, 83, [87](#), 94, 102, 107, 319
PMIX_NODE_INFO_ARRAY, 314, 318, [323](#), 323, 327, 328, 333, 334, 600
PMIX_NODE_LIST, [88](#), 90, 92, 551
PMIX_NODE_MAP, [88](#), 90, 92, 315, 326, 327, 341, 342, 589
PMIX_NODE_MAP_RAW, [88](#), 602
PMIX_NODE_OVERSUBSCRIBED, [95](#), 318, 619
PMIX_NODE_RANK, [93](#), 320, 454
PMIX_NODE_SIZE, [95](#), 318
PMIX_NODEID, 80, 83, [87](#), 90, [94](#), 94, 95, 102, 105, 107, 110, 112, 275, 286, 289, 314, 318, 320,

323, 334, 551, 600, 607
 PMIX_NOHUP, 435, 440, **443**, 602
 PMIX_NOTIFY_COMPLETION, **190**, 435, 457, 557, 563, 567
 PMIX_NOTIFY_JOB_EVENTS, **190**, 435, 456, 609
 PMIX_NOTIFY_PROC_ABNORMAL_TERMINATION, **190**, 609
 PMIX_NOTIFY_PROC_TERMINATION, **190**, 609
 PMIX_NPROC_OFFSET, **89**, 316
 PMIX_NSDIR, **90**, 94, 319, 320
 PMIX_NSPACE, 87, 92, **93**, 103–105, 108–111, 113, 313, 315, 322, 323, 327, 394, 395, 452, 453,
 457, 460, 558, 563, 566, 598–600
 PMIX_NUM_ALLOCATED_NODES, **88**, 601
 PMIX_NUM_NODES, 73, 86, **88**, 89, 91, 326, 327, 551, 601
 PMIX_NUM_SLOTS, **88**, 89, 92
 PMIX_OPTIONAL, 77, 79, 82, **84**
 PMIX_OUTPUT_TO_DIRECTORY, **189**, 609
 PMIX_OUTPUT_TO_FILE, 178, 184, **189**, 380
 PMIX_PACKAGE_RANK, **93**, 320, 601
 PMIX_PARENT_ID, **93**, 378, 441, 617
 PMIX_PERSISTENCE, 127–132, **133**, 371, 481
 PMIX_PERSONALITY, 177, 183, **188**, 379
 PMIX_PPR, 177, 183, **188**, 379
 PMIX_PREFIX, 176, 182, **188**, 378
 PMIX_PRELOAD_BIN, 177, 183, **188**, 379
 PMIX_PRELOAD_FILES, 177, 183, **188**, 379
 PMIX_PREPEND_ENVAR, 179, 185, **191**, 569
 PMIX_PRIMARY_SERVER, **432**, 466, 602
 PMIX_PROC_INFO, **87**, 92, 103, 107
 PMIX_PROC_INFO_ARRAY, 313, 319, **322**, 328, 600, 613
 PMIX_PROC_MAP, **88**, 89, 90, 92, 316, 326, 327, 341, 342, 551, 589
 PMIX_PROC_MAP_RAW, **88**, 602
 PMIX_PROC_PID, **94**, 105, 110
 PMIX_PROC_STATE_STATUS, 105, 110, **458**
 PMIX_PROC_TERM_STATUS, 457, **458**
 PMIX_PROCDIR, **94**, 320
 PMIX_PROCID, 92, **93**, 103, 105, 107–109, 111, 190, 191, 313, 314, 323, 395, 435–437, 456, 557,
 563, 567, 600, 610
 PMIX_PROGRAMMING_MODEL, 66, **68**, 317, 342, 571
 PMIX_PSET_MEMBERS, 246, **247**, 607
 PMIX_PSET_NAME, 246, **247**, 607
 PMIX_PSET_NAMES, 245, **247**, 317, 579, 608
 PMIX_QUERY_ALLOC_STATUS, 105, 110, **113**, 395
 PMIX_QUERY_ATTRIBUTE_SUPPORT, 103, 104, 108, **111**, 118, 452, 598
 PMIX_QUERY_AUTHORIZATIONS, 105, 110, **111**

PMIX_QUERY_AVAIL_SERVERS, [112](#), 429, 598
 PMIX_QUERY_DEBUG_SUPPORT, 104, 109, [111](#), 394, 556
 PMIX_QUERY_GROUP_MEMBERSHIP, [250](#), 579, 608
 PMIX_QUERY_GROUP_NAMES, [250](#), 579, 608
 PMIX_QUERY_JOB_STATUS, 104, 109, [111](#), 394
 PMIX_QUERY_LOCAL_ONLY, [113](#), 395
 PMIX_QUERY_LOCAL_PROC_TABLE, 104, 109, 111, 394, 452, [460](#), 558, 563, 564
 PMIX_QUERY_MEMORY_USAGE, 104, 109, [111](#), 395
 PMIX_QUERY_NAMESPACE_INFO, [110](#), 598
 PMIX_QUERY_NAMESPACES, 104, 109, [110](#), 394, 453, 567
 PMIX_QUERY_NUM_GROUPS, [250](#), 579, 608
 PMIX_QUERY_NUM_PSETS, 112, [247](#), 607
 PMIX_QUERY_PROC_TABLE, 104, 109, 111, 394, 452, [460](#), 563, 564, 566
 PMIX_QUERY_PSET_MEMBERSHIP, 112, [247](#), 607
 PMIX_QUERY_PSET_NAMES, 112, [247](#), 607
 PMIX_QUERY_QUALIFIERS, 106, [112](#), 112, 598
 PMIX_QUERY_QUEUE_LIST, 104, 109, [111](#), 394
 PMIX_QUERY_QUEUE_STATUS, 104, 109, [111](#), 394
 PMIX_QUERY_REFRESH_CACHE, 102, 106, 107, [112](#), 117
 PMIX_QUERY_REPORT_AVG, 105, 109, [113](#), 395
 PMIX_QUERY_REPORT_MINMAX, 105, 109, [113](#), 395
 PMIX_QUERY_RESULTS, 105, [112](#), 598
 PMIX_QUERY_SPAWN_SUPPORT, 104, 109, [111](#), 394, 556
 PMIX_QUERY_STORAGE_LIST, [476](#), 616
 PMIX_QUERY_SUPPORTED_KEYS, [110](#), 598
 PMIX_QUERY_SUPPORTED_QUALIFIERS, [110](#), 598
 PMIX_RANGE, 127, 129, 130, 132, [133](#), 136–138, 143, 145, 150, 233, 250, 371, 374, 376, 391,
 421, 422, 481, 579, 608
 PMIX_RANK, [93](#), 103, 105, 108, 109, 111, 180, 186, 313, 314, 319, 323, 395, 455, 460, 552, 557,
 561, 600
 PMIX_RANKBY, 177, 183, [188](#), 316, 379
 PMIX_REGISTER_CLEANUP, 225, 228, [230](#)
 PMIX_REGISTER_CLEANUP_DIR, 225, 228, [231](#)
 PMIX_REGISTER_NODATA, 313, [322](#)
 PMIX_REINCARNATION, [94](#), 320, 602
 PMIX_REPORT_BINDINGS, 178, 184, [189](#), 381
 PMIX_REQUESTOR_IS_CLIENT, 378, [382](#)
 PMIX_REQUESTOR_IS_TOOL, 378, [382](#), 567
 PMIX_REQUIRED_KEY, 369, [370](#), 601
 PMIX_RM_NAME, [88](#), 314
 PMIX_RM_VERSION, [88](#), 314
 PMIX_SEND_HEARTBEAT, 233, 235, [236](#)
 PMIX_SERVER_ATTRIBUTES, 103, 108, [113](#), 118, 592, 599

PMIX_SERVER_ENABLE_MONITORING, 307, [309](#)
PMIX_SERVER_FUNCTIONS, 103, 104, 109, 111, [113](#), 118, 598, 599
PMIX_SERVER_GATEWAY, 305, [309](#)
PMIX_SERVER_HOSTNAME, 315, [431](#)
PMIX_SERVER_INFO_ARRAY, 112, [113](#), 598
PMIX_SERVER_NAMESPACE, 305, [309](#), 315, 429, 461, 465
PMIX_SERVER_PIDINFO, 429, [431](#), 461, 465
PMIX_SERVER_RANK, 305, [309](#), 315
PMIX_SERVER_REMOTE_CONNECTIONS, 306, [309](#)
PMIX_SERVER_SCHEDULER, 287, 290, 305, [310](#), 599, 604
PMIX_SERVER_SESSION_SUPPORT, 305, [309](#), 599
PMIX_SERVER_SHARE_TOPOLOGY, 307, [309](#), 599
PMIX_SERVER_START_TIME, [309](#), 599
PMIX_SERVER_SYSTEM_SUPPORT, 305, [309](#), 426
PMIX_SERVER_TMPDIR, 305, 307, [309](#), 426–428
PMIX_SERVER_TOOL_SUPPORT, 296, 305, 308, [309](#)
PMIX_SERVER_URI, 105, 110, 428, 429, [431](#), 461, 465
PMIX_SESSION_ID, 87, 89, [93](#), 96, 313, 314, 322, 326, 457, 600
PMIX_SESSION_INFO, 80, 83, [86](#), 88, 96, 102, 107, 314, 315, 342
PMIX_SESSION_INFO_ARRAY, 313, 314, [322](#), 323, 326, 587
PMIX_SET_ENVAR, 179, 185, [191](#), 569
PMIX_SET_SESSION_CWD, 176, 182, [189](#), 378
PMIX_SETUP_APP_ALL, 340, [343](#)
PMIX_SETUP_APP_ENVARS, 340, [343](#), 557, 562
PMIX_SETUP_APP_NONENVARS, 340, [343](#)
PMIX_SINGLE_LISTENER, 65, 306, [309](#)
PMIX_SINGLETON, 307, [310](#), 620
PMIX_SOCKET_MODE, 65, 306, [309](#), 462
PMIX_SPAWN_TIMEOUT, 180, 186, [191](#), 381, 620
PMIX_SPAWN_TOOL, 180, 186, [190](#), 442, 561
PMIX_SPAWNED, [94](#), 320, 378
PMIX_STDIN_TGT, 177, 184, [189](#), 379
PMIX_STORAGE_ACCESS_TYPE, [477](#), 617
PMIX_STORAGE_ACCESSIBILITY, [476](#), 616
PMIX_STORAGE_BW_CUR, [477](#), 477, 617
PMIX_STORAGE_BW_MAX, [477](#), 616
PMIX_STORAGE_CAPACITY_LIMIT, [476](#), 616
PMIX_STORAGE_CAPACITY_USED, [476](#), 616
PMIX_STORAGE_ID, [476](#), 476, 616
PMIX_STORAGE_IOPS_CUR, [477](#), 477, 617
PMIX_STORAGE_IOPS_MAX, [477](#), 617
PMIX_STORAGE_MEDIUM, [476](#), 616
PMIX_STORAGE_MINIMAL_XFER_SIZE, [477](#), 616

PMIX_STORAGE_OBJECT_LIMIT, [477](#), [616](#)
PMIX_STORAGE_OBJECTS_USED, [477](#), [616](#)
PMIX_STORAGE_PATH, [476](#), [616](#)
PMIX_STORAGE_PERSISTENCE, [476](#), [616](#)
PMIX_STORAGE_SUGGESTED_XFER_SIZE, [477](#), [477](#), [616](#), [617](#)
PMIX_STORAGE_TYPE, [476](#), [616](#)
PMIX_STORAGE_VERSION, [476](#), [616](#)
PMIX_SWITCH_PEERS, [290](#), [290](#), [605](#)
PMIX_SYSTEM_TMPDIR, [305](#), [309](#), [426](#), [428](#)
PMIX_TAG_OUTPUT, [177](#), [184](#), [189](#), [380](#)
PMIX_TCP_DISABLE_IPV4, [65](#), [68](#), [306](#), [462](#)
PMIX_TCP_DISABLE_IPV6, [65](#), [68](#), [306](#), [463](#)
PMIX_TCP_IF_EXCLUDE, [65](#), [68](#), [306](#), [462](#)
PMIX_TCP_IF_INCLUDE, [65](#), [68](#), [306](#), [462](#)
PMIX_TCP_IPV4_PORT, [65](#), [68](#), [306](#), [462](#)
PMIX_TCP_IPV6_PORT, [65](#), [68](#), [306](#), [462](#)
PMIX_TCP_REPORT_URI, [65](#), [67](#), [306](#), [462](#)
PMIX_TCP_URI, [68](#), [428](#), [429](#), [461](#), [465](#)
PMIX_TDIR_RMCLEAN, [88](#), [316](#)
PMIX_THREADING_MODEL, [66](#), [68](#), [571](#)
PMIX_TIME_REMAINING, [100](#), [105](#), [110](#), [111](#), [395](#)
PMIX_TIMEOUT, [4](#), [77](#), [78](#), [81](#), [84](#), [85](#), [96](#), [122](#), [124](#), [127](#), [129](#), [130](#), [132](#), [136](#), [138](#), [143](#), [145](#), [180](#),
[186](#), [191](#), [196](#), [199](#), [200](#), [202](#), [249](#), [253](#), [255](#), [258](#), [259](#), [261](#), [262](#), [266](#), [267](#), [269](#), [298](#), [299](#),
[301](#), [303](#), [366](#), [370](#), [372](#), [374](#), [376](#), [381](#), [383](#), [386](#), [410](#), [413](#), [432](#), [468](#), [563](#), [604](#), [620](#)
PMIX_TIMEOUT_REPORT_STATE, [190](#), [609](#)
PMIX_TIMEOUT_STACKTRACES, [190](#), [609](#)
PMIX_TIMESTAMP_OUTPUT, [178](#), [184](#), [189](#), [380](#)
PMIX_TMPDIR, [87](#), [88](#), [90](#), [316](#), [318](#), [319](#)
PMIX_TOOL_ATTACHMENT_FILE, [428](#), [429](#), [432](#), [461](#), [465](#), [602](#)
PMIX_TOOL_ATTRIBUTES, [103](#), [108](#), [113](#), [118](#), [592](#), [599](#)
PMIX_TOOL_CONNECT_OPTIONAL, [432](#), [602](#)
PMIX_TOOL_DO_NOT_CONNECT, [428](#), [432](#), [461](#), [463](#)
PMIX_TOOL_FUNCTIONS, [103](#), [104](#), [109](#), [111](#), [113](#), [118](#), [598](#), [599](#)
PMIX_TOOL_NAMESPACE, [396](#), [427](#), [430](#), [461](#), [463](#)
PMIX_TOOL_RANK, [396](#), [427](#), [431](#), [461](#), [463](#)
PMIX_TOPOLOGY2, [307](#), [309](#), [599](#), [612](#)
PMIX_UNIV_SIZE, [9](#), [87](#), [314](#), [326](#), [586](#), [587](#), [589](#)
PMIX_UNSET_ENVAR, [179](#), [185](#), [191](#), [569](#)
PMIX_USERID, [127](#), [129](#), [130](#), [132](#), [136](#), [138](#), [143](#), [144](#), [216](#), [219](#), [224](#), [227](#), [232](#), [234](#), [238](#), [241](#),
[297](#), [299](#), [300](#), [302](#), [371–374](#), [376](#), [378](#), [388](#), [394](#), [396](#), [398](#), [399](#), [401](#), [405](#), [408](#), [410](#), [412](#),
[414](#), [415](#), [419](#)
PMIX_USOCK_DISABLE, [65](#), [306](#), [309](#)
PMIX_VERSION_INFO, [397](#), [398](#)

PMIX_WAIT, 83, [85](#), 136–138, 374
PMIX_WAIT_FOR_CONNECTION, [432](#), 468, 563, 604
PMIX_WDIR, 176, 182, [188](#), 317, 378

PMIX_ALLOC_NETWORK

Deprecated, [612](#)

PMIX_ALLOC_NETWORK_ENDPTS

Deprecated, [613](#)

PMIX_ALLOC_NETWORK_ENDPTS_NODE

Deprecated, [613](#)

PMIX_ALLOC_NETWORK_ID

Deprecated, [612](#)

PMIX_ALLOC_NETWORK_PLANE

Deprecated, [612](#)

PMIX_ALLOC_NETWORK_QOS

Deprecated, [612](#)

PMIX_ALLOC_NETWORK_SEC_KEY

Deprecated, [613](#)

PMIX_ALLOC_NETWORK_TYPE

Deprecated, [612](#)

PMIX_ARCH

Deprecated, [591](#)

Removed, [614](#)

PMIX_COLLECTIVE_ALGO

Deprecated, [591](#)

Removed, [614](#)

PMIX_COLLECTIVE_ALGO_REQD

Deprecated, [588](#)

Removed, [614](#)

PMIX_DEBUG_JOB

Deprecated, [612](#)

PMIX_DEBUG_WAIT_FOR_NOTIFY

Deprecated, [617](#)

PMIX_DSTPATH

Deprecated, [591](#)

Removed, [613](#)

PMIX_ERROR_GROUP_ABORT

Deprecated, [586](#)

Removed, [588](#)

PMIX_ERROR_GROUP_COMM

Deprecated, [586](#)

Removed, [588](#)

PMIX_ERROR_GROUP_GENERAL

Deprecated, [586](#)

Removed, [588](#)
PMIX_ERROR_GROUP_LOCAL
Deprecated, [586](#)
Removed, [588](#)
PMIX_ERROR_GROUP_MIGRATE
Deprecated, [586](#)
Removed, [588](#)
PMIX_ERROR_GROUP_NODE
Deprecated, [586](#)
Removed, [588](#)
PMIX_ERROR_GROUP_RESOURCE
Deprecated, [586](#)
Removed, [588](#)
PMIX_ERROR_GROUP_SPAWN
Deprecated, [586](#)
Removed, [588](#)
PMIX_ERROR_HANDLER_ID
Deprecated, [586](#)
Removed, [588](#)
PMIX_ERROR_NAME
Deprecated, [586](#)
Removed, [588](#)
PMIX_HWLOC_HOLE_KIND
Deprecated, [591](#)
Removed, [613](#)
PMIX_HWLOC_SHARE_TOPO
Deprecated, [591](#)
Removed, [613](#)
PMIX_HWLOC_SHMEM_ADDR
Deprecated, [591](#)
Removed, [613](#)
PMIX_HWLOC_SHMEM_FILE
Deprecated, [591](#)
Removed, [613](#)
PMIX_HWLOC_SHMEM_SIZE
Deprecated, [591](#)
Removed, [613](#)
PMIX_HWLOC_XML_V1
Deprecated, [591](#)
Removed, [613](#)
PMIX_HWLOC_XML_V2
Deprecated, [591](#)
Removed, [613](#)

PMIX_LOCAL_TOPO
 Deprecated, [591](#)
 Removed, [613](#)

PMIX_LOCALITY
 Deprecated, [613](#)

PMIX_MAP_BLOB
 Deprecated, [591](#)
 Removed, [614](#)

PMIX_MAPPER
 Deprecated, [591](#)
 Removed, [614](#)

PMIX_NON_PMI
 Deprecated, [591](#)
 Removed, [614](#)

PMIX_PROC_BLOB
 Deprecated, [591](#)
 Removed, [614](#)

PMIX_PROC_DATA
 Deprecated, [613](#)

PMIX_PROC_URI
 Deprecated, [591](#)
 Removed, [614](#)

PMIX_RECONNECT_SERVER
 Deprecated, [612](#)

PMIX_TOPOLOGY
 Deprecated, [612](#)

PMIX_TOPOLOGY_FILE
 Deprecated, [591](#)
 Removed, [613](#)

PMIX_TOPOLOGY_SIGNATURE
 Deprecated, [592](#)
 Removed, [613](#)

PMIX_TOPOLOGY_XML
 Deprecated, [592](#)
 Removed, [613](#)