Reflections Of The Void

A blog about life, Engineering, Business, Research, and everything else (especially everything else)

Wednesday, July 08, 2020

Software RDMA revisited : setting up SoftiWARP on Ubuntu 20.04

Almost ten years ago I wrote about installing Softlwarp on Ubuntu 10.04. Today I will be revisiting the process. First, what is Softlwarp: Soft-iWARP is a software-based iWARP stack that runs at reasonable performance levels and seamlessly fits into the OFA RDMA environment provides several benefits. SoftiWARP is a software RDMA device that attaches with the active network cards to enable RDMA programming. For anyone starting with RDMA programming, RDMA-enabled hardware might not be at hand. SoftiWARP is a very useful tool to set up the RDMA environment, and code and experiments with.

To install Softlwarp you have to go through 4 stages: Setting up the environment, Building Softlwarp, Configuring Softiwarp, Testing.

Setting up RDMA environment

Before you start you should prepare the environment for building a kernel module and userspace library.
Basic building environment

```
sudo apt-get install build-essential libelf-dev cmake
```

Installing userspace libraries and tools

```
sudo apt-get install libibverbs1 libibverbs-dev librdmacm1 \
librdmacm-dev rdmacm-utils ibverbs-utils
```

Insert common RDMA kernel modules

```
sudo modprobe ib_core
sudo modprobe rdma_ucm
```

Check if everything is correctly installed :

```
sudo lsmod | grep rdma
```

You should see something like this :

```
rdma_ucm               28672  0
ib_uverbs             126976  1 rdma_ucm
rdma_cm                61440  1 rdma_ucm
iw_cm                  49152  1 rdma_cm
ib_cm                  57344  1 rdma_cm
ib_core               311296  5 rdma_cm,iw_cm,rdma_ucm,ib_uverbs,ib_cm
```

Now set up some library for the userspace libs :

```
sudo apt-get install build-essential cmake gcc libudev-dev libnl-3-dev \
libnl-route-3-dev ninja-build pkg-config valgrind
```

Installing SoftiWARP

10 years ago you had to clone the SoftiWARP source code and build it (https://github.com/zrlio/softiwarp.git). Now you are lucky, it is by default in the Linux kernel 5.3 and above!

You just have to type :

```
sudo modprobe siw
```

verify it works :

```
sudo lsmod | grep siw
```

you should see :

Reflections:

Essays

Short Posts

Featured post

Yet another Red Queen Project : Franco-German Gaia-X

For some reason, the EU and especially the French government love moonshot project. The only problem is that they tend to be launched after ...

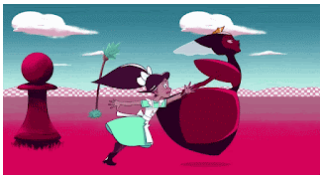

Popular Posts

[Links of the Day] 15/01/2019 : Incident Response best practice, Database Schema Crawler, Fingerprinting TLS

ja3 : something I discovered recently. Apparently, you can fingerprint SSL and TLS session in order to identify the service being run behi...

How to install Soft-iWARP on Ubuntu 10.10, AKA how to have RDMA enabled system without the expensive hardware.

What is Soft-iWARP and why installing it : Soft-iWARP [1] is a software-based iWARP stack that runs at reasonable performance...

Software RDMA revisited : setting up SoftiWARP on Ubuntu 20.04

Almost ten years ago I wrote about installing Softlwarp on Ubuntu 10.04. Today I will be revisiting the process. First, what is Softlwarp: ...

Wake up call, cloud is not green.

In 2008 I mentioned multiple times the issue with virtualization , cloud computing and the false impression that it help save the planet....

[Links of the Day] 08/11/2018 : large scale study of datacenter network reliability, What to measure in production, Failure Mode effect Analysis

A Large Scale Study of Data Center Network Reliability : the authors study reliability within and between Facebook datacenters. One of the...

Search This Blog

Search

```
siw                   188416  0
ib_core               311296  6 rdma_cm,iw_cm,rdma_ucm,ib_uverbs,siw,ib_cm
libcrc32c              16384  3 nf_conntrack,nf_nat,siw
```

moreover, you should check if you have an Infiniband device present :

```
ls /dev/infiniband
```

Result :

```
rdma_cm
```

You also need to add the following file in your /etc/udev/rules.d/90-ib.rules directory containing the below entries :

```
####  /etc/udev/rules.d/90-ib.rules  ####
KERNEL=="umad*", NAME="infiniband/%k"
KERNEL=="issm*", NAME="infiniband/%k"
KERNEL=="ucm*", NAME="infiniband/%k", MODE="0666"
KERNEL=="uverbs*", NAME="infiniband/%k", MODE="0666"
KERNEL=="uat", NAME="infiniband/%k", MODE="0666"
KERNEL=="ucma", NAME="infiniband/%k", MODE="0666"
KERNEL=="rdma_cm", NAME="infiniband/%k", MODE="0666"
########
```

If it doesn't exist you need to create it.

I would suggest you add also the module to the list of modules to load at boot by adding them to /etc/modules file

You need now to reboot your system.

Userspace library

Normally, recent library support softiwarp out of the box. But if you want to compile your own version follow the step bellow. However, do this at your own risk... I recommend to stick with the std libs.

Optional build SIW userland libraries:

All the userspace library are in a nice single repository. You just have to clone the repo and build all the shared libraries. If you want you can also just build **libsiw** but it's just easier to build everything at once.

```
git clone https://github.com/zrlio/softiwarp-user-for-linux-rdma.git
cd ./softiwarp-user-for-linux-rdma/
./buid.sh
```

Now we have to setup the $LD_LIBRARY_PATH so that build libraries can be found.

```
cd ./softiwarp-user-for-linux-rdma/build/lib/
export LD_LIBRARY_PATH=$(pwd):$LD_LIBRARY_PATH
```

or you can add the line in your .bashrc profile:

```
export LD_LIBRARY_PATH=<<PATHTOTHELIBRARIES>>:$LD_LIBRARY_PATH
```

End of optional section

Setup the SIW interface :

Now we will be setting up the loopback and a standard eth interface as RDMA device:

```
sudo rdma link add <NAME OF SIW DEVICE > type siw netdev <NAME OF THE INTERFACE>
```

In this case for me :

```
sudo rdma link add siw0 type siw netdev enp0s31f6
sudo rdma link add siw_loop type siw netdev l0
```

You can check the two devices have been correctly set up using **ivc_devices** and **ibv_devinfo** command result of **ibv_devices** :

```
    device                 node GUID
    ------              ----------------
    siw0                507b9ddd7a170000
    siw_loop            0000000000000000
```

result of **ibv_devinfo** :

```
hca_id: siw0
 transport:   iWARP (1)
 fw_ver:    0.0.0
 node_guid:   507b:9ddd:7a17:0000
 sys_image_guid:   507b:9ddd:7a17:0000
 vendor_id:   0x626d74
 vendor_part_id:   0
 hw_ver:    0x0
 phys_port_cnt:   1
  port: 1
   state:   PORT_ACTIVE (4)
   max_mtu:  1024 (3)
   active_mtu:  invalid MTU (0)
   sm_lid:   0
   port_lid:  0
   port_lmc:  0x00
   link_layer:  Ethernet
hca_id: siw_loop
 transport:   iWARP (1)
 fw_ver:    0.0.0
 node_guid:   0000:0000:0000:0000
 sys_image_guid:   0000:0000:0000:0000
 vendor_id:   0x626d74
 vendor_part_id:   0
 hw_ver:    0x0
 phys_port_cnt:   1
  port: 1
   state:   PORT_ACTIVE (4)
   max_mtu:  4096 (5)
   active_mtu:  invalid MTU (0)
   sm_lid:   0
   port_lid:  0
   port_lmc:  0x00
   link_layer:  Ethernet
```

Testing with RPING:

Now we simply test the setup with rping :

In one shell :

```
rping -s -a <serverIP>
```

in the other :

```
rping -c -a <serverIP> -v
```

And you should see the rping working successfully!

You are now all set to use RDMA without the need for expensive hardware.

Posted by Benoit Hudzia at 2:26 pm
Labels: rdma , Softlwarp , ubuntu

No comments :

Post a Comment

Newer Post Home Older Post

Total Pageviews

Disclaimer