INSTALLING INTEL MPSS 3.3 IN ARCH LINUX

Andrey Vladimirov
Colfax International

August 20, 2014

Abstract

This technical publication provides instructions for installing the Intel Manycore Platform Software Stack (MPSS) version 3.3 in Arch Linux operating system. Intel MPSS is a suite of tools necessary for operation of Intel Xeon Phi coprocessors. Instructions provided here enable offload and networking functionality for coprocessors in Arch Linux. The procedure described in this paper is completely reversible via an uninstallation script.

1. SUPPORTED DISTRIBUTIONS FOR INTEL MPSS

Intel MPSS [2] is a suite of tools and drivers necessary for the operation of Intel Xeon Phi coprocessors. As of the writing of this paper, MPSS is supported for the following operating systems:

- Red Hat* Enterprise Linux* 64-bit 6.2 with Linux kernel 2.6.32-220
- Red Hat* Enterprise Linux* 64-bit 6.3 with Linux kernel 2.6.32-279
- Red Hat* Enterprise Linux* 64-bit 6.4 with Linux kernel 2.6.32-358
- Red Hat* Enterprise Linux* 64-bit 6.5 with Linux kernel 2.6.32-431
- Red Hat* Enterprise Linux* 64-bit 7.0 with Linux kernel 3.10.0-123
- SUSE* Linux* Enterprise Server 11 SP2 with kernel 3.0.13-0.27-default
- SUSE* Linux* Enterprise Server 11 SP3 with kernel 3.0.76-0.11-default
- Microsoft Windows* 7 Enterprise SP1 (64-bit)
- Microsoft Windows* 8/8.1 Enterprise (64-bit)
- Microsoft Windows* Server 2008 R2 SP1 (64-bit)
- Microsoft Windows* Server 2012 (64-bit)
- Microsoft Windows* Server 2012 R2 (64-bit)

If an operating system is not supported, it does not mean that MPSS will not work; it only means Intel developers have not validated MPSS for that system. For example, CentOS Linux with versions corresponding to the respective RHEL versions allow MPSS to work "out of the box".

This publication reports a method of enabling Intel MPSS on Arch Linux. According to DistroWatch.com [3], Arch Linux was one of the top 10 most popular Linux distributions in 2014.

These instructions apply only to MPSS 3.3. Other versions of MPSS may work as well, however, they have not been tested. In the tested configuration, we verified offload and network connectivity with coprocessors.

2. DOWNLOADS

You can rely on `wget` to obtain the necessary software. Instructions for that are incorporated in the instructions. Alternatively, prior to installing MPSS on Arch Linux, download the following software:

1) Intel MPSS 3.3 archive, `mpss-3.3-linux.tar`, via direct link or from [2]. If version 3.3 is not available via direct link or on the MPSS page, it may still be available in the MPSS archive [4].
2) Tarball of long-term support Linux kernel 3.10, `linux-lts310.tar.gz`. This archive is available from AUR [5].
3) TRIG utility `trig.sh` from Colfax Research [1], and
4) A set of RHEL networking scripts `rhnet.tgz` from the same page [1].

The GNU General Public License, version 2 (GPLv2) applies to TRIG and RHEL networking scripts package. The terms of this license are available at [6].

3. PREREQUISITES

Prior to starting this installation procedure, be aware that we have made the best effort we could to make these instructions correct and useful. However, we do not imply warranty or any kind, and assume no liability for any general, special, incidental or consequential damages arising out of the use of these instructions or software employed in this procedure.

These instructions assume that the Arch Linux system has working Internet connection, or is connected to an Arch Linux repository mirror. All operations must be performed under the `root` account.

Optionally, you can set the environment variable `PS1` as shown below so that the command prompt reflects the current directory. This will make it easier to compare your terminal output to the listings shown in this paper, and to control your location in the directory tree.

```
[root@lyra ~]# export PS1='[\u@\h \w]\$ ' # do not copy-paste this command, type
[root@lyra ~]# cd /usr/bin
[root@lyra /usr/bin]# # <-- Your location is now clear
```

Begin by installing the following prerequisite packages:

```
[root@lyra ~]# pacman -S base base-devel
[root@lyra ~]# pacman -S openssh wget rpmextract python bc net-tools
```

Then place the LTS kernel tarball from AUR at `/root/linux-lts310.tar.gz` and install it:

```
[root@lyra ~]# cd /root
[root@lyra ~]# wget http://uni.colfax-intl.com/s/linux-lts310.tar.gz # redirect
[root@lyra ~]# tar -xf linux-lts310.tar.gz
[root@lyra ~/linux-lts310]# cd linux-lts310
[root@lyra ~/linux-lts310]# makepkg --asroot -s
[root@lyra ~/linux-lts310]# pacman -U linux-lts310-3.10.53-1-x86_64.pkg.tar.xz
[root@lyra ~/linux-lts310]# pacman -U linux-lts310-headers-3.10.53-1-*.tar.xz
```

In order to activate the LTS kernel, re-generate the GRUB configuration file using `grub-mkconfig` and reboot the system:

```
[root@lyra ~]# grub-mkconfig -o /boot/grub/grub.cfg
[root@lyra ~]# reboot
```

After reboot, confirm that "`uname -r`" reports using the LTS kernel 3.10:

```
[root@lyra ~]# uname -r
```

The method of booting the LTS kernel may be different if you are using a bootloader other than GRUB. More information can be found in the ArchWiki [7].

4. PREPARING THE INSTALLATION TREE AND SCRIPTS

If you are relying on `wget` to download the required software, continue to follow these instructions. Otherwise (i.e., if you had downloaded the software manually), ignore the calls to `wget` in the instructions, and place the MPSS archive at `/root/mpss-3.3-linux.tar`, place the RHEL network scripts at `/root/rhnet.tgz`, and the TRIG utility at `/root/trig.sh`.

Obtain and unpack the MPSS archive and other software and create the Arch installation tree location:

```
[root@lyra ~]# cd /root/
[root@lyra ~]# wget http://uni.colfax-intl.com/s/mpss-3.3-linux.tar #redirect
[root@lyra ~]# wget http://uni.colfax-intl.com/s/rhnet.tgz # redirect
[root@lyra ~]# wget http://uni.colfax-intl.com/s/trig.sh #redirect
[root@lyra ~]# chmod +x trig.sh
[root@lyra ~]# tar -xf mpss-3.3-linux.tar
[root@lyra ~]# mkdir /root/mpss-3.3/Arch
[root@lyra ~]# cd /root/mpss-3.3/Arch/
[root@lyra ~/mpss-3.3/Arch]#
```

Unpack the RHEL network tools. On top of them, unpack the MPSS RPMs:

```
[root@lyra ~/mpss-3.3/Arch]# tar -xf /root/rhnet.tgz
[root@lyra ~/mpss-3.3/Arch]# for x in ../*rpm; do rpmextract.sh $x; done
[root@lyra ~/mpss-3.3/Arch]# ls # View the installation tree
bin  etc  opt  sbin  usr
```

Now several steps must be performed to configure the installation tree for Arch Linux:

```
[root@lyra ~/mpss-3.3/Arch]# mkdir -p etc/systemd/system
[root@lyra ~/mpss-3.3/Arch]# mv etc/mpss/mpss.service etc/systemd/system/
[root@lyra ~/mpss-3.3/Arch]# mkdir -p etc/init.d
[root@lyra ~/mpss-3.3/Arch]# mv etc/mpss/mpss.redhat etc/init.d/mpss
[root@lyra ~/mpss-3.3/Arch]# mkdir -p etc/prelink.conf.d
[root@lyra ~/mpss-3.3/Arch]# echo \
>   -b /usr/bin/miccheck > etc/prelink.conf.d/miccheck.conf
[root@lyra ~/mpss-3.3/Arch]# mv etc/mpss/micras.redhat etc/init.d/micras
[root@lyra ~/mpss-3.3/Arch]# mkdir -p usr/lib64/python3.4/site-packages
[root@lyra ~/mpss-3.3/Arch]# ln -s \
>   /usr/src/miccheck/_miccheck usr/lib64/python3.4/site-packages/_miccheck
[root@lyra ~/mpss-3.3/Arch]# ln -s \
>   /usr/src/micmgmt usr/lib64/python3.4/site-packages/micmgmt
[root@lyra ~/mpss-3.3/Arch]#
```

Use the TRIG script to generate installation and uninstallation scripts for the tree built in `/root/mpss-3.3/Arch`:

```
[root@lyra ~/mpss-3.3]# /root/trig.sh /root/mpss-3.3/Arch/
Creating tree-install-Arch.sh to install the tree of files in /root/mpss-3.3/Arch
Creating tree-uninstall-Arch.sh to remove files installed by tree-install-Arch.sh
[root@lyra ~/mpss-3.3]#
```

5. COMPILING THE MIC MODULE

Before installing MPSS, compile, install and load the MIC kernel module:

```
[root@lyra ~/mpss-3.3]# mkdir /root/mpss-3.3/rebuild
[root@lyra ~/mpss-3.3]# cd /root/mpss-3.3/rebuild/
[root@lyra ~/mpss-3.3/rebuild]# rpmextract.sh ../src/mpss-modules-3.3-1.src.rpm
[root@lyra ~/mpss-3.3/rebuild]# tar -xf mpss-modules-3.3.tar.bz2
[root@lyra ~/mpss-3.3/rebuild]# make MIC_CARD_ARCH=k1om &> log-make.txt
[root@lyra ~/mpss-3.3/rebuild]# make MIC_CARD_ARCH=k1om install &> log-install.txt
[root@lyra ~/mpss-3.3/rebuild]# depmod
[root@lyra ~/mpss-3.3/rebuild]# modprobe mic
[root@lyra ~/mpss-3.3/rebuild]# lsmod | grep "mic "
mic                    639834  0
```

6. INSTALLING AND STARTING MPSS

Now install MPSS and create initial configuration using `micctrl`:

```
[root@lyra ~]# /root/mpss-3.3/tree-install-Arch.sh # Put all files in place
[root@lyra ~]# ssh-keygen # Do this for all users
[root@lyra ~]# micctrl --initdefaults # Initialize MPSS configuration
```

Start MPSS and enable MPSS loading on system boot. Also, enable service `micnet`. This is a custom add-on included in `rhnet.tgz` required for enabling MIC networking at boot

```
[root@lyra ~]# systemctl enable mpss
[root@lyra ~]# systemctl start mpss
[root@lyra ~]# systemctl enable micnet
```

7. TESTING FUNCTIONALITY

In order to confirm that the coprocessors are booted and accessible via SSH, run

```
[root@lyra ~]# ssh mic0
[root@lyra-mic0 ~]# echo I have \
>   `cat /proc/cpuinfo | grep processor | wc -l` hardware threads
I have 240 hardware threads
[root@lyra-mic0 ~]# exit
```

If you have already installed the Intel compiler, download and run a test offload application:

```
[root@lyra ~]# source /opt/intel/composerxe/bin/compilervars.sh intel64
[root@lyra ~]# wget http://uni.colfax-intl.com/s/Offload-Hello.cc #redirect
[root@lyra ~]# icpc -o Offload-Hello Offload-Hello.cc
[root@lyra ~]# ./Offload-Hello
Hello World from main()! I see 32 logical cores.
Hello World from offload! I see 240 logical cores.
```

8. STOPPING AND UNINSTALLING MPSS

In order to uninstall MPSS, the following steps must be performed:

```
[root@lyra ~]# systemctl stop mpss
[root@lyra ~]# systemctl disable mpss
[root@lyra ~]# systemctl disable micnet
[root@lyra ~]# micctrl --cleanconfig
[root@lyra ~]# rmmod mic
[root@lyra ~]# /root/mpss-3.3/tree-uninstall-Arch.sh
[root@lyra ~]# # Optionally, also remove the MIC module:
[root@lyra ~]# rm -f /lib/modules/3.10.53-1-lts310/extra/mic.ko
```

ACKNOWLEDGEMENTS

Thanks to Dr. Christopher Lupo of California Polytechnic State University for requesting and testing this installation procedure.

NEWSLETTER

Receive our Newsletter to be notified about new Colfax Research publications, educational materials and news on parallel programming. Subscribe at http://nlreg.colfax-intl.com/. It is completely free, and you can unsubscribe any time.

REFERENCES

[1] Landing page for this paper, "Installing Intel MPSS 3.3 in Arch Linux". http://research.colfaxinternational.com/post/2014/08/20/Arch.aspx.

[2] Intel Manycore Platform Software Stack (MPSS). http://software.intel.com/en-us/articles/intel-manycore-platform-software-stack-mpss.

[3] DistroWatch. http://distrowatch.com/dwres.php?resource=popularity.

[4] Intel Manycore Platform Software Stack (MPSS) Archive. http://software.intel.com/en-us/articles/intel-manycore-platform-software-stack-mpss-archive.

[5] Linux LTS 3.10 kernel in Arch User Repository (AUR). https://aur.archlinux.org/packages/linux-lts310/.

[6] GNU General Public License v2.0. http://www.gnu.org/licenses/gpl-2.0.html.

[7] Enhance system stability - ArchWiki. https://wiki.archlinux.org/index.php/Enhance_system_stability#Install_the_linux-lts_package.