

ADINA DMP System 9.3

Installation Notes

for

Linux (only)

Updated for version 9.3.2

**ADINA R & D, Inc.
71 Elton Avenue
Watertown, MA 02472**

**support@adina.com
www.adina.com**

Table of contents

1. About ADINA DMP System 9.3	3
2. ADINA DMP System prerequisites	3
2.1 Open MPI 1.4.3	4
2.1.1 Open MPI 1.4.3 built and distributed by ADINA R&D	4
2.1.2 Open MPI 1.4.3 built and installed at your site	4
3. Installing ADINA DMP	5
3.1 How to obtain a DMP floating license	5
3.2 Mounting the CD-ROM drive	5
3.3 Selecting the ADINA DMP System installation directory	6
3.4 Preparing to set up ADINA DMP	6
3.5 Installing ADINA DMP	6
4. Customizing the Open MPI parameters and environment variables used by ADINA DMP	8
4.1 Open MPI MCA parameters	8
4.2 PSM environment variables	9
5. Troubleshooting the installation	9
6. Third-party software distributed with the ADINA DMP distribution	9
7. Configurations used by ADINA R&D for building and testing ADINA DMP	10
7.1 2-node cluster	10
7.2 4-node cluster	11
7.3 16-node computer	12
8. Support for ADINA DMP	12
8.1 Support before ADINA DMP is installed	12
8.2 Support after ADINA DMP is installed	13

1. About ADINA DMP System 9.3

The ADINA Distributed Memory Processing (DMP) System 9.3 is available for clusters composed of x86_64 computers running Linux (kernel 2.6 or higher). ADINA DMP can also be used on a single computer with multi-core CPUs.

Currently, ADINA DMP is capable of solving ADINA, ADINA-F and ADINA-FSI models.

For general information about ADINA DMP System 9.3, see the *ADINA DMP System 9.3 User Guide*. The Glossary at the end of that document should be referred to for definitions of computing terms used in ADINA DMP.

In order to use the ADINA DMP System, a special DMP floating license is required. Please see the next sections for the complete list of hardware and software prerequisites, including how to obtain a DMP floating license.

There are many different possible hardware and software combinations in a cluster, and ADINA R&D cannot support all possible hardware and software combinations in a cluster. Section 7 of this document describes the hardware and software for which ADINA R&D has tested ADINA DMP. For other hardware and software combinations, ADINA R&D will provide a “reasonable” amount of support. See Section 8 of this document for instructions for obtaining support.

2. ADINA DMP System prerequisites

The following are the minimum hardware and software prerequisites for the ADINA DMP System:

- One or more x86_64 computers running Linux with kernel 2.6 or higher.
- A special DMP floating license, which must be obtained from ADINA R & D, Inc. (see below).
- Open MPI 1.4.3 (see below)

If ADINA DMP is to be used in a cluster, then there are additional prerequisites:

- the communication protocols rsh and rlogin (and ssh if chosen during installation of ADINA DMP) on the cluster must work without a password requirement,
- the firewall must be turned off on all nodes of the cluster,
- the hypervisor capability in Red Hat Enterprise Linux 7 and later versions must be disabled—the interface virbr0 must be shut down and the service libvirtd disabled,

- the ADINA DMP System must be installed in a shared directory (e.g., a directory in a file system mounted on all the cluster nodes via NFS)

2.1 Open MPI 1.4.3

You can use the distribution of Open MPI 1.4.3 built and distributed by ADINA R&D, or you can use the distribution of Open MPI 1.4.3 built and installed at your site. **It is strongly recommended that you use the distribution of Open MPI 1.4.3 built and installed at your site, because the distribution from ADINA R&D might not contain support for the networking at your site.**

2.1.1 Open MPI 1.4.3 built and distributed by ADINA R&D.

The ADINA DMP System distribution includes shared libraries and executables from Open MPI 1.4.3. Therefore it is not required that Open MPI 1.4.3 be installed at your site. The Open MPI distribution supplied by ADINA R&D includes some support for InfiniBand, however we cannot guarantee that the networking at your site is supported by this Open MPI distribution. For example, the Open MPI distribution supplied by ADINA R&D does not include support for InfiniPath.

Open MPI 1.4.3 supplied by ADINA R&D is built on the 4-node cluster configuration described in Section 7.2.

In the Open MPI distribution supplied by ADINA R&D for versions 9.3.1 and earlier, there was support for the Torque PBS scheduler (Torque version 2.4.16). However the current Open MPI distribution supplied by ADINA R&D does not include support for the Torque PBS scheduler. If you use Torque at your site, you should use Open MPI installed at your site, see instructions below.

If you choose to install our distribution of Open MPI, then, after installation of ADINA DMP, you can use the command `<ADINA_DMP_HOME>/ompi_info_wrapper` to obtain information about our Open MPI distribution.

See Section 6 for more information about the third-party software distributed by ADINA R&D.

2.1.2 Open MPI 1.4.3 built and installed at your site

If you do not have Open MPI 1.4.3 installed at your site, you can download the Open MPI 1.4.3 source from www.open-mpi.org. Please consult your system administrator for assistance in downloading, building and installing Open MPI 1.4.3.

You can build Open MPI 1.4.3 using either the GNU compilers or the Intel compilers.

The installation of Open MPI 1.4.3 must be in a shared directory, that is, in a directory that can be accessed from any node in the cluster.

Remember, if you update your networking hardware, you should rebuild Open MPI. Open MPI does not auto-detect changes in networking hardware.

3. Installing ADINA DMP

3.1 How to obtain a DMP floating license

A license server machine must have been set up previously. This is done with the standard (non-DMP) ADINA System installation disk. Note that the license server machine does not necessarily need to be a node in the cluster. It can be any external machine that is accessible to the cluster via a network.

On the machine to be used as the ADINA floating license server,

1. Run the `adr1mid` command (in `<ADINAHOME>/r1m`) to obtain the hostid for the server.

```
cd <ADINAHOME>/r1m
./adr1mid
```

Send the output of the above `adr1mid` command to ADINA R & D (password@adina.com).

2. Once you receive the floating license file (e.g., `yourcomp.lic`), put the file in the `<ADINAHOME>/r1m` directory and start the `rlm` license server program using the command `rlm`.
3. `./rlm -dlog r1m.log &`

where `r1m.log` is the server debug log file.

For more information on the RLM License Administration, please refer to the document `RLM_License_Administration.pdf` in `<ADINAHOME>/r1m`. Note that not all options described in the document are available in the ADINA System.

3.2 Mounting the CD-ROM Drive

If you are installing ADINA System 9.3 from the self extracting file downloaded from the ADINA ftp site, please skip to Section 3.3.

Insert the ADINA DMP System 9.3 CD in the CD-ROM drive and mount the drive if it is not automatically mounted.

The typical mount command in Linux is

```
mount /mnt/cdrom (or /media/cdrom)
```

Please consult your system administrator if you have a problem mounting the CD-ROM drive.

3.3 Selecting the ADINA DMP System installation directory

Select the directory to install the ADINA DMP System, i.e., the ADINA DMP System installation directory or <ADINA_DMP_HOME>. Note that the directory must be a shared directory on your cluster. Create the directory if necessary, then change the working directory to the ADINA DMP System installation directory.

For example, if the ADINA DMP System installation directory is /usr/adina_dmp93, the commands are

```
umask 022
mkdir /usr/adina_dmp93
cd /usr/adina_dmp93
```

3.4 Preparing to set up ADINA DMP

If you are installing from the CD, copy the contents of the CD to <ADINA_DMP_HOME>.

If you are installing from the self-extracting file addmp9.3.#.tar.gz, place the file in the <ADINA_DMP_HOME> directory, unzip and untar it using the command

```
tar xzvf addmp9.3.#.tar.gz
```

3.5 Installing ADINA DMP

Run the setup_dmp script to install ADINA DMP.

```
cd <ADINA_DMP_HOME>
chmod +x setup_dmp
./setup_dmp
```

During the setup you will be asked for the following information:

- Directory containing the standard ADINA distribution.
- Type of ADINA floating license that you have.
- Name or IP address of the floating license server.
- Whether to use the Open MPI distribution provided by ADINA R&D, or whether to use an Open MPI distribution at your site.
- If using an Open MPI distribution at your site, the directory containing the Open MPI distribution to use. This Open MPI distribution must be an Open MPI 1.4.3 distribution.
- Remote execution tool that ADINA DMP will use on your cluster (either rsh or ssh).

- For each BTL component of Open MPI, whether ADINA DMP should use that BTL component. It is recommended that you allow ADINA DMP to use each BTL component.
- For each MTL component of Open MPI, whether ADINA DMP should use that MTL component. It is recommended that you allow ADINA DMP to use each MTL component. (Note, your installation might not have any MTL components of Open MPI.)

If, after installation, you decide that you want to change the above information, you can run `setup_dmp` again to make the changes.

BTL and MTL components of Open MPI: message transports

Open MPI uses the BTL and MTL components to control the message transport of MPI messages between processes. Thus, by choosing which BTL and MTL components ADINA DMP can use, you can control the speed of message passing between processes.

For example, on the 2-node cluster at ADINA R&D (see Section 7), the following BTL components are available:

self: used for loopback (`setup_dmp` always automatically selects the self component)
sm: shared memory, used for message passing between processes in a single node
tcp: message passing using an Ethernet network
openib: message passing using an InfiniBand network
ofud: obsolete (`setup_dmp` never selects the ofud component)

and the following MTL components are available:

psm: message passing using Performance Scaled Messaging using an InfiniPath network

If you use your own distribution of Open MPI, you might not see all of these components, and you might see other components.

When ADINA DMP is invoked, Open MPI automatically chooses the components to use for message passing, from the components that are selected using `setup_dmp`.

Notice that Open MPI might choose between several possible components. For example, if your hardware supports InfiniPath, you select all possible components with `setup_dmp`, and if all processes are running on the local node, Open MPI can choose the MTL `psm` component, the BTL `openib` component and/or the BTL `sm` component. Open MPI attempts to choose the components corresponding to the fastest way to pass MPI messages. This is why we recommend that you allow ADINA DMP to use all components.

However, you might find that Open MPI does not always make the best choice for your particular cluster. Then you can use `setup_dmp` to choose only the components of your choice.

4. Customizing the Open MPI parameters and environment variables used by ADINA DMP

The information given in this section is intended for the advanced system administrator who needs to modify the way in which ADINA DMP uses Open MPI. For example, the cluster might consist of nodes connected using a networking protocol not considered by our setup script.

For installations similar to those given in Section 7, it should not be necessary to do any of the customizations in this section.

All of the customizations described in this section are performed by editing the `addmp9.3` script.

4.1 Open MPI MCA parameters

The Open MPI configuration process is quite complex and there are many tuneable parameters associated with Open MPI. These tuneable parameters are referred to as MCA parameters. For general information about the MCA parameters, see the documentation at www.open-mpi.org.

By default (that is, when the `setup_dmp` utility is used to install ADINA DMP), the script `addmp9.3` is edited so that ADINA DMP uses the following MCA parameters:

`--mca plm_rsh_agent`: Set to either `rsh:ssh` or `ssh:rsh`, depending on the remote executable tool chosen.

`--mca orte_tmpdir_base`: Set to `/tmp`. It is assumed that `/tmp` is a local directory (not a shared directory) for each node in the cluster.

`--mca btl_openib_warn_default_gid_prefix`: Set to 0. This setting suppresses a warning message that is printed for certain InfiniBand networks.

`--mca btl`: Either blank, or set to a list of BTL components to include. This list is built from the BTL components chosen when the `setup_dmp` script is run. For example, supposing that BTL `tcp` component is chosen and the BTL `openib` component is not chosen, the corresponding MCA parameter string is `--mca btl self,tcp`.

`--mca mt1`: Either blank, or set to a list of MTL components to include. This list is built from the MTL components chosen when the `setup_dmp` script is run. For example, supposing that MTL `psm` component is chosen, the corresponding MCA parameter string is `--mca mt1 psm`.

Exception, if there is more than one MTL component available and no MTL components are chosen, the corresponding MCA parameter string is `--mca mt1 ^(list of MTL components)`. For example, supposing that only the MTL `psm` component is available, and that the `psm` component is not chosen, the corresponding MCA parameter string is `--mca mt1 ^psm`.

You can find all of these parameter settings in the `addmp9.3` script. Search for “`set mca`”.

4.2 PSM environment variables

When the MTL `psm` component is used, there are additional environment variables that influence the message transport using PSM. One important environment variable is `IPATH_NO_CPUAFFINITY`. `IPATH_NO_CPUAFFINITY` should be unset if ADINA DMP is used in pure DMP mode (that is, when the number of threads is set to 1), and `IPATH_NO_CPUAFFINITY` should be set to 1 if ADINA DMP is used in hybrid mode (that is, when the number of threads is greater than 1). The `addmp9.3` script automatically sets `IPATH_NO_CPUAFFINITY` based on these conditions.

5. Troubleshooting the installation

To troubleshoot the installation, see the *ADINA DMP System 9.3 User Guide*, Section 9 “Troubleshooting ADINA DMP”. We have put the troubleshooting information into that document because the basic concepts of DMP, as well as the basic usage of ADINA DMP, are explained in that document. This basic information is necessary in order to troubleshoot the installation.

6. Third-party software distributed with the ADINA DMP distribution

The following third-party software is distributed with the ADINA DMP distribution:

Open MPI 1.4.3: `<ADINA_DMP_HOME>/openmpi-1.4.3`

Built from the tarball `openmpi-1.4.3.tar.gz`, downloaded from `www.open-mpi.org`. The following build options are explicitly set when running the `configure` script:

```
--prefix=/opt/openmpi-1.4.3-adinabuilddist
CC=gcc
CXX=g++
F77=gfortran
FC=gfortran
--disable-mca-dso
--with-openib
--enable-mca-no-build=ras-tm,plm-tm
```

libibverbs-1.1.4: `<ADINA_DMP_HOME>/libibverbs-1.1.4`

librdmacm-1.0.13: `<ADINA_DMP_HOME>/librdmacm-1.10.13`

libgfortran-4.1.2: `<ADINA_DMP_HOME>/libgfortran-4.1.2`

See these directories for license information.

libibverbs-1.1.4, librdmacm-1.0.13 and libgfortran-4.1.2 are distributed because the distribution of Open MPI 1.4.3 supplied by ADINA R&D requires them. Libibverbs and librdmacm are from the MLNX_OFED distribution for InfiniBand, version MLNX_OFED_LINUX-1.5.2-1.0.0.

7. Configurations used by ADINA R&D for building and testing ADINA DMP

This section gives the specifications for some of the configurations used by ADINA R&D for building and testing ADINA DMP. All of these configurations are located at the ADINA R&D main office in Watertown, MA.

7.1 2-node cluster

ADINA R&D tests ADINA DMP on a 2-node cluster with the following configuration:

Number of nodes	2
Number of CPUs per node	1
Number of cores per node	4
Total number of cores	8
CPU information	Node 0: Intel Xeon E31225, processor speed 3.1 GHz Node 1: AMD PRO A10-8850B R7, processor speed 1.7 GHz
Hyper-threading	disabled
Memory (RAM)	31 GB for node 0, 30 GB for node 1
Operating system	Centos 3.10.0-327.13.1.el7
Fortran compiler used to build Open MPI	gfortran version 4.8.5
C compiler used to build Open MPI	gcc version 4.8.5
InfiniBand adapter for each node	Intel QLE-7342
IB stack	OFED-3.18-2
Open MPI	1.4.3
Open MPI BTL components available	self, sm, openib, tcp
Open MPI MTL components available	psm

7.2 Four-node cluster

ADINA R&D builds and tests ADINA DMP on a four-node cluster with the following configuration:

Number of nodes	4
Number of CPUs per node	2
Number of cores per node	8
Total number of cores	32
CPU information	Intel Xeon E5640, processor speed 2.67 GHz
Hyper-threading	disabled
Memory (RAM)	48 GB for node 0, 24 GB for the other nodes
Operating system	Red Hat Enterprise Linux Server release 5.5, kernel 2.6.18-195el5
Fortran compiler used to build ADINA DMP	Intel ifort version 11.1
C compiler used to build ADINA DMP	Intel icc version 11.1
Fortran compiler used to build Open MPI	gfortran version 4.1.2
C compiler used to build Open MPI	gcc version 4.1.2
Interconnect for Ethernet	HP GbE2C Layer 2/3 Ethernet Blade Switch
InfiniBand adapter for each node	Mellanox MT26428 (ConnectX VPI), 2 ports, firmware version 2.8.0
InfiniBand switch	Mellanox Infiniscale-IV
IB stack	MLNX_OFED_LINUX-1.5.2-1.0.0
Open MPI	1.4.3
Open MPI BTL components available	self, sm, tcp, openib, ofud
Open MPI MTL components available	none

7.3 16-core computer

ADINA R&D runs ADINA DMP on a single 16-core computer with the following configuration:

Number of nodes	1
Number of CPUs per node	2
Number of cores per node	8
Total number of cores	16
CPU information	Intel Xeon E5-2687W, processor speed 3.1 GHz
Hyper-threading	disabled
Memory (RAM)	256 GB
Operating system	Red Hat Enterprise Linux Server release 6.2, kernel 2.6.32-250el6
Fortran compiler used to build Open MPI	Intel ifort version 11.1
C compiler used to build Open MPI	Intel icc version 11.1
Open MPI	1.4.3
Open MPI BTL components available	self, sm, tcp
Open MPI MTL components available	none

8. Support for ADINA DMP

If you need additional support for ADINA DMP, please contact ADINA R&D.

In order to provide support, ADINA R&D requires information about your computer or cluster, similar to the information given in the tables in Section 7. Please send us the output from the commands discussed below.

Please note that ADINA R&D does not provide support for the setup of networking between nodes in a cluster. We assume that, before you contact ADINA support, you have successfully networked the nodes into a cluster.

8.1 Support before ADINA DMP is installed

Information about Open MPI

Please send us the output from the command

```
ompi_info -a
```

For example, supposing that you have installed Open MPI 1.4.3 in the /opt/openmpi-1.4.3 directory, the command is

```
/opt/openmpi-1.4.3/bin/ompi_info -a
```

Information about InfiniBand

If your cluster uses InfiniBand, the command

```
ofed_info
```

provides much useful information about the InfiniBand IB stack. Other InfiniBand commands at your installation can be determined with the command

```
man -k infiniband
```

Some InfiniBand commands that might be present at your installation, and that give useful information for support, are

```
iblinkinfo
```

Report link info for all links in the fabric

```
ibnodes
```

Show InfiniBand nodes in topology

```
ibnetdiscover
```

Discover InfiniBand topology

Note that you might need root privileges in order to run these commands.

8.2 Support after ADINA DMP is installed

See Section 9 of the *ADINA DMP System 9.3 User Guide* for some troubleshooting commands. The output from these troubleshooting commands is very useful, however this output can be difficult to interpret. ADINA R&D can assist in interpreting this output. Using `openmpiprobe9.3` with the `-mca info all` option gives the most output.