

mental ray Installation Notes for UNIX

Document Version 3.0
September 24, 2004

Copyright Information

Copyright © 1986-2005 mental images GmbH, Berlin, Germany.

All rights reserved.

This document is protected under copyright law. The contents of this document may not be translated, copied or duplicated in any form, in whole or in part, without the express written permission of mental images GmbH.

The information contained in this document is subject to change without notice. mental images GmbH and its employees shall not be responsible for incidental or consequential damages resulting from the use of this material or liable for technical or editorial omissions made herein.

mental images®, incremental images™, mental ray®, mental matter®, mental ray Phenomenon®, mental ray Phenomena™, Phenomenon™, Phenomena™, Phenomenon Creator™, Phenomenon Editor™, Photon Map™, mental ray Relay™ Library, Relay™ Library, SPM®, Shape-By-Shading™, Internet Rendering Platform™, iRP™, Reality™, RealityServer®, RealityPlayer™, RealityDesigner™, iray®, imatter®, neuray™, mental mill™, MetaSL™, Meta Shading Language™, Network Rendering Manager™, and NRM™ are trademarks or, in some countries, registered trademarks of mental images GmbH, Berlin, Germany.

All other product names mentioned in this document may be trademarks or registered trademarks of their respective companies and are hereby acknowledged.

1 Introduction

This document describes installation of mental ray. For the installation of the SPM licensing software look in the document **SPM Installation Manual**.

For an easy installation of mental ray the installation script `rayinstall` is available. If you want to do the installation manually or need some hints for troubleshooting follow the instructions in this manual.

NOTE

We strongly recommend that installation will be performed by the systems administrator, as it requires superuser status and involves changing of system files.

Installation is straightforward with the exception of licensing and network parallelism. For the former, a network daemon (the Software Protection Manager) must be installed. For the latter, a network daemon is required on every machine which is to be used as a rendering server.

The Software Protection Manager (SPM) protects mental ray from unauthorized use. While executing, mental ray communicates with the SPM daemon to verify authorization. Since communication may take place over the network, mental ray may run on a host other than the one which runs the SPM. Thus, the software may be copied for backup or may run on any host which has access, via the network, to the SPM.

Note that when running mental ray in the network mode, one license is required for each CPU mental ray runs on. These licenses can be obtained from a single or from multiple SPM daemon(s), and these may be located anywhere in the network.

Before you begin the installation, you have to decide on which machines you will run the Software Protection Manager daemons, that is, on which machines you will store your licenses. An SPM host should be fast (the SPM daemon should be able to respond quickly to the protected software, although it does not consume much CPU time). A heavily loaded NFS server is a particular bad SPM server because the NFS server code in the kernel can prevent the SPM daemon from responding to requests in time. We strongly discourage you to choose a machine that is directly connected to the Internet, because the SPM License Server is not designed to resist port attacks on your machine. So it may open an unwanted root access into your machine. Currently we are not aware of any existing exploits.

`mental ray` can be run in network mode. That means you can run it as client and server application. Clients will be invoked from the command line and servers from `inetd` or `xinetd` as a service.

The smallest configuration will use one machine with an SPM daemon and a ray client; on larger sites, you will have a bunch of machines both as client and server, and some of them running also the SPM daemon.

2 Software Installation

The first step in installation is mounting the CD on `/CDROM`. CD's are in Rock Ridge and Joliet format. We recommend to put mental ray into `/usr/local/mi` (which is allowed to be a soft link to somewhere else). A list of supported platforms can be found at the `mental images` website www.mentalimages.com.

If you want to use the mental ray installation script simply type:

```
rayinstall
```

in the mental ray distribution directory.

Calling `rayinstall` will display the installation menu. Choosing 'full installation' will install mental ray and the software protection manager.

Invoking `rayinstall` without any option allows the user to customize the configuration (customize mode). A lot of informational messages will be displayed. Normally, it is recommended to use this customize mode.

```
rayinstall -d perform a default installation
rayinstall -s no information messages are displayed
rayinstall -u deinstall mental ray
rayinstall -h get help
```

The default installation directory for mental ray is `/usr/local/mi` but it may be changed in customize mode. Executables will be copied to `/usr/local/mi/bin`. All header files, that are necessary for mental ray will be copied to `/usr/local/mi/rayinc`. The shader library will be installed in `/usr/local/mi/shaders`.

In default mode (`rayinstall -d`) the SPM Daemon will be installed on the local host. In customize mode, you have the option of installing the SPM Daemon on a remote host if it is possible to login there without a password. `rayinstall` will also prepare the local host for network rendering. In customize mode you can do this for remote hosts too. The appropriate changes in `/etc/services`, `/etc/hosts` and `/etc/inetd.conf` (`/etc/xinetd.d`) are automatically done. Original files will be saved with the extension `orig`.

Note: If your site is running NIS for the distribution of services or hosts the common NIS database will not be changed. You have to do it manually. If DNS is used for name resolution the host database will not be changed.

`rayinstall` will start the SPM Daemon and change the boot scripts. If you do not want `rayinstall` to do this for you, run `rayinstall` in customize mode.

Users who decide not to use `rayinstall` or who need a detailed explanation of what exactly `rayinstall` is doing, should install mental ray manually as described below.

Guidelines for manual installation:

```
<become superuser>
mount the CD to /CDROM
mkdir /usr/local/mi
cp -r /CDROM/PLATFORM/bin /usr/local/mi
cp -r /CDROM/PLATFORM/shaders /usr/local/mi
cp -r /CDROM/common/include /usr/local/mi/rayinc
cp -r /CDROM/PLATFORM/spm /var
```

PLATFORM describes your special UNIX platform. The SPM utilities should go to `/var/spm` in any case. Note that `/var/spm` may not be NFS-mounted because the SPM programs, running as root, have limited access to NFS partitions.

mental ray Client Installation

1. `mental ray` needs to know where in the network the SPM License Server can be found. It will automatically search for a free SPM license on the machines nicknamed `spmhost0`, `spmhost1`..., up to `spmhost9`. Ask your network administrator about name resolution in your network (local hosts, NIS or DNS) and which `hostdatabase` has to be changed.

A local entry in `/etc/hosts` should look like:

```
172.15.0.0    castor    spmhost0
172.15.0.1    pollux    spmhost1
```

Do this on every machine you want to run `mental ray` on.

If your site is running DNS (the domain name service), the `hosts` name server tables must be updated instead of modifying all `/etc/hosts`. If your site is running NIS, the `hosts` NIS map must be updated instead of modifying all `/etc/hosts`.

Alternatively any user can set the environment variable `SPM_HOST` in his private environment to a preferred SPM License Servers's hostname.

In a C-shell environment type:

```
setenv SPM_HOST castor
```

in a Bourne- or Korn-shell environment:

```
SPM_HOST=castor
export SPM_HOST
```

`mental ray` will try to get it's license first from `castor` and if this was not successful from `spmhost0`, `spmhost1`....

`SPM_HOST` can also be set to more than one SPM License Server.

2. In order to find shaders `mental ray` will look for a file `rayrc` using the following search path:

```
current directory
home directory
/usr/local/mi
```

In the current directory or home directory the file has to be called `.rayrc` (with a leading dot). If `mental ray` was not installed at the default location `/usr/local/mi` the environment variable `MI_ROOT` has to be set to the installation directory.

The syntax of this `rayrc` file is `mi Syntax`. An example for a `rayrc` file could be:

```
link "/usr/local/mi/shaders/softimage.so"
```

3. Execution of `mental ray` should be possible at this point if licenses have been already installed..

4. To use dynamic linking for your own user-defined shaders, you need a C compiler and linker which are unbundled from the operating system by some computer vendors. However, shaders distributed in DSO form by a shader vendor can be used without having a C compiler.

For shader compilation the file `shader.h` must be publicly accessible on each `ray` client and server. We recommend to put it into

```
/usr/local/mi/rayinc
```

As a minimal configuration, the program `ray` may be run by itself with an appropriately formatted input file. Typically, however, `ray` is called from a client application. Also, one or more translator programs, such as `sdltoimi` or `wftoimi` will be used to input data from other systems. The utility `imf_disp` may be used for displaying images and the utility `imf_copy` is used to convert image files to other formats including the format required to memory-map texture images into `ray`.

mental ray Server Installation

If `mental ray` is to be run concurrently over several loosely coupled systems (i.e., over FDDI or Ethernet), it must be called from the `inetd` as a network server.

In addition, the presence of a `.rayhosts` file is required. For further details, refer to *mental ray User's Manual and Reference*.

Installation of the `ray` server requires superuser status and is accomplished as follows:

1. `mental ray` servers communicate with the client (master) and from server to server via a TCP port (`mi-ray`). A Internet service number has to be assigned. Add a line like:

```
mi-ray      7002/tcp      # mental images ray tracer
```

to `/etc/services`.

If another service uses number 7002, any number is fine, but make sure that that this entry is consistent on all render servers.

If your site is running NIS (formerly called Yellow Pages) and NIS is used for services, ask your network administrator to change the NIS service database.

2. The `mental ray` network server will be started by the `inetd` or `xinetd` whenever a request comes in on port `mi-ray`. If your system runs `inetd` the configuration file `/etc/inetd.conf` has to be edited. Add the following line:

```
mi-ray  stream  tcp  nowait  nobody  /usr/local/mi/bin/ray  ray
```

Systems running `xinetd`, need a new file entry in `/etc/xinetd.d` called `mi-ray`. This file `mi-ray` contains the following lines:

```
service mi-ray
{
    disable           = no
    socket_type       = stream
    wait              = no
    user              = nobody
    server            = /usr/local/mi/bin/ray
    groups            = yes
    flags             = REUSE
}
```

WARNING

The `ray` server opens a security hole by letting the user specify the command name of the C compiler for dynamic linking; a malicious user could call a program to change passwords or destroy data. Thus we strongly recommend that the `ray` server be run with a `userid` allowed nothing but creating files in `/tmp` (e.g. `nobody`), but if you don't care about security, the `ray` server could also be run as `root`.

Note: Due to a problem in SGI and HP systems software, the traditional user and/or group ID of -2 for nobody will cause the operating system to fail in its attempt to invoke the ray server. We recommend changing these IDs to any unique positive value less than 32768.

3. The inetd has to be refreshed after changing its configuration file. Use `ps` to look for inetd's process id, and send it a HUP signal to force it to reread its configuration file.

For example, on a generic System V machine

```
ps -ef | fgrep inetd
root 115      1  0  Dec 20  ?           0:04 /usr/etc/inetd
kill -HUP 115
```

SGI and Linux:

```
killall -HUP inetd
```

AIX:

```
refresh -s inetd
```

HP-UX:

```
inetd -c
```

4. To test the installation, create a `.rayhosts` file in the current directory with the name(s) of the machine(s) of the planned render servers (one name per line).

Now try to execute `ray` with the `-verbose` option. If everything has been set up correctly, `ray` will tell you that the render servers (e.g. `castor` and `pollux`) are connected. You should see the following output:

```
MSG 0.0 info :  reading hosts file ../rayhosts
MSG 0.0 info :  connecting host castor:7002
MSG 0.0 info :  connecting host pollux:7002
```

If e.g. `castor` is not configured as a render server you will see the following message:

```
MSG 0.0 info :  reading hosts file ../rayhosts
MSG 0.0 error 011301:  cannot connect host castor:7002
MSG 0.0 info :  connecting host pollux:7002
```

An error message like:

```
MSG 0.0 info :  reading hosts file /your/homedir/.rayhosts
MSG 0.0 info :  adding new host 1 (castor:7002)
MSG 0.0 info :  adding new host 2 (pollux:7002)
MSG 0.0 error:  cannot receive welcome message from host 2 (active)
MSG 0.0 error:  cannot add host pollux:7002
```

indicates that *castor* is configured correctly but *pollux* failed to connect.

For more information, use `ray -verbose 7` to print more detailed debugging messages. Press Control-C or Control-D to exit `ray` after this test.

If the loading fails, or `ray` is not even able to connect to the remote server, look at the remote host's syslog file for inetd's error messages. If it fails after connection was set up additional error messages can be found in `/tmp/raylib.log`.