

Service Pack Roll: Users Guide



Version 5.4.2 Edition



Service Pack Roll: Users Guide :

Version 5.4.2 Edition

Published Jan 06 2011

Copyright © 2011 University of California

This document is subject to the Rocks License (see Appendix A: Rocks Copyright).

Table of Contents

Preface	v
1. Overview	1
2. Installing	3
2.1. On a New Server	3
2.2. On an Existing Server.....	3
2.3. Adding the Roll to Nodes.....	4
3. Fixes Included in the Service Pack Roll	5
A. Rocks Copyright	7

List of Tables

1-1. Summary..... 1
1-2. Roll Compatibility 1

Preface

The Rocks Service Pack Roll contains fixes for Rocks version 5.4.2.

Chapter 1. Overview

Table 1-1. Summary

Name	service-pack
Version	5.4.2
Maintained By	Rocks Group
Architecture	i386, x86_64
Compatible with Rocks®	5.4.2

Table 1-2. Roll Compatibility


Roll	Requires ^a	Optional ^b	Conflicts
alpha		X	
area51		X	
base	X		
bio		X	
condor		X	
ganglia		X	
grid		X	
hpc		X	
java		X	
kernel	X		
os (disk 1)		X	
os (disk 2)		X	
os (disk 3)		X	
os (disk 4)		X	
os (disk 5)		X	
os (disk 6)		X	
os (disk 7)		X	
pbs		X	
service-pack	X		
sge		X	
viz		X	
web-server		X	
xen		X	

Roll	Requires ^a	Optional ^b	Conflicts
<p>Notes:</p> <p>a. You may also substitute your own OS CDs for the Rocks® OS Roll CDs. In this case you must use all the CDs from your distribution and not use any of the Rocks® OS Roll CDs.</p> <p>b. Only Rolls that have been verified as compatible with this Roll are listed. Other Rolls will likely work, but have not been tested by the maintainer of this Roll.</p>			

Chapter 2. Installing

2.1. On a New Server

The service-pack Roll should be installed during the initial installation of your server (or cluster). This procedure is documented in section 1.2 of the Rocks® usersguide. You should select the service-pack Roll from the list of available rolls when you see a screen that is similar to the one below.



2.2. On an Existing Server

The Service Pack Roll can be installed on running frontend.

The following procedure will install and configure the roll for the frontend.

For i386, download the Service Pack Roll:

```
# wget http://www.rocksclusters.org/ftp-site/pub/rocks/rocks-5.4/linux/service-pack-5.4.2-1.i386.d
```

Or for x86_64, download the Service Pack Roll:

```
# wget http://www.rocksclusters.org/ftp-site/pub/rocks/rocks-5.4/linux/service-pack-5.4.2-1.x86_64.d
```


Then, as root, make sure to remove any old Service Pack Rolls:

```
# rocks remove roll service-pack
```

Now apply the new Service Pack Roll:

```
# rocks set host attr localhost roll_install_on_the_fly true shadow=yes
# rocks add roll service-pack*iso
# rocks enable roll service-pack
# cd /export/rocks/install
# rocks create distro
# rocks run roll service-pack | bash
# init 6
```

2.3. Adding the Roll to Nodes

To add the the Service Pack Roll to the nodes of your cluster, you'll need to reinstall the nodes.

To reinstall all your nodes, as root execute:

```
# rocks run host '/boot/kickstart/cluster-kickstart-pxe'
```



It is critical that you run `cluster-kickstart-pxe` as it will force the compute nodes to PXE boot. It is important that you PXE boot the nodes for the first install, because with a PXE boot based install, the nodes will get their `initrd` from the frontend and inside the `initrd` is a new `tracker-client` that is compatible with the new `tracker-server`. After the first install, a new `initrd` will be on the hard disk of the installed nodes and then it is safe to run `/boot/kickstart/cluster-kickstart`.

Chapter 3. Fixes Included in the Service Pack Roll

Base Roll

- In specific configurations, when you run "rocks sync users", you could see the error message:

```
411-alert: error while loading shared libraries: librocks.so: cannot open shared
object file: No such file or directory
```

This occurs when 'ldconfig' is not run during the post installation. The fix is to ensure ldconfig is run at the end of an installation for all nodes.

- Fixes to the Rocks Command Line and 411-master program to fully support channel bonding on the frontend.
- When only one alias was specified for a host (with the command: `rocks add host alias`), the entry in `/var/named/private.domain` would have the format similar to `compute-1-0 CNAME ('c1-0',)`.
The fix to the Rocks Command Line now properly outputs `compute-1-0 CNAME c1-0`.
- Clients were unable to use the Rocks Command Line. This was because a file in `/etc/ld.so.conf.d` was not installed on the clients that allowed mysql client code to find its libraries.
Clients can now run Rocks Command Line commands.
- The ability to build a "login appliance" was only available when the SGE Roll was supplied.
The login appliance configuration logic has been moved into the Base Roll, so now one can build login appliances on any Rocks clusters.
- Memory leak fixed for gmond.
- Properly report disk partitions into the database for software RAID file systems.
Prior to this fix, software RAID file systems on client nodes would be reformatted on subsequent installations.
- Increase installation speeds for clients with software RAID file systems.

Prior to this fix, if one of the partitions is a software RAID and if the raid is in a "dirty" state, mdadm will try to 'resync' the disks to ensure all the data is protected. On a running system, this is a great thing to do. On an installing system, it slows the install down considerably.

We didn't find a way to disable the 'resync' action for mdadm, but we can extremely throttle back the top rebuild speed by executing:

```
echo 1 > /proc/sys/dev/raid/speed_limit_min
echo 2 > /proc/sys/dev/raid/speed_limit_max
```

Kernel Roll

- There was a code in rocks-tracker that wasn't properly reconstructing the 'peers' circular list when that list was expanded. This could cause several MD5 checksum errors on installing nodes because an installing node would ask for a peer and the tracker would send stale peer data to the installing node.
- In the case where one had a 64-bit frontend and 32-bit compute nodes, the messages exchanged between the tracker-server and the tracker-clients on installing nodes would be misaligned and installing nodes were incapable of finding (and thus, incapable of downloading) their packages.

All tracker messages are now properly aligned between 32-bit and 64-bit machines.

Appendix A. Rocks Copyright

Rocks(r)
www.rocksclusters.org
version 5.4 (Maverick)

Copyright (c) 2000 - 2010 The Regents of the University of California.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice unmodified and in its entirety, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising and press materials, printed or electronic, mentioning features or use of this software must display the following acknowledgement:

"This product includes software developed by the Rocks(r)
Cluster Group at the San Diego Supercomputer Center at the
University of California, San Diego and its contributors."

4. Except as permitted for the purposes of acknowledgment in paragraph 3, neither the name or logo of this software nor the names of its authors may be used to endorse or promote products derived from this software without specific prior written permission. The name of the software includes the following terms, and any derivatives thereof: "Rocks", "Rocks Clusters", and "Avalanche Installer". For licensing of the associated name, interested parties should contact Technology Transfer & Intellectual Property Services, University of California, San Diego, 9500 Gilman Drive, Mail Code 0910, La Jolla, CA 92093-0910, Ph: (858) 534-5815, FAX: (858) 534-7345, E-MAIL:invent@ucsd.edu

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

