

# **Service Pack Roll: Users Guide**



**Version 5.0.3 Edition**



**Service Pack Roll: Users Guide :**

Version 5.0.3 Edition

Published Jul 2008

Copyright © 2008 University of California

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

# Table of Contents

<b>Preface</b> .....	<b>v</b>
<b>1. Overview</b> .....	<b>1</b>
<b>2. Installing</b> .....	<b>2</b>
2.1. On a New Server .....	2
2.2. Adding the Roll to a Running System.....	2
<b>3. Fixes Included in the Service Pack Roll</b> .....	<b>4</b>
<b>A. Rocks Copyright</b> .....	<b>5</b>

# List of Tables

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

# Preface

The Rocks Service Pack Roll contains fixes for Rocks version 5.0.3.

# Chapter 1. Overview

**Table 1-1. Summary**

Name	service-pack
Version	5.0.3
Maintained By	Rocks Group
Architecture	i386, x86_64
Compatible with Rocks™	5.0.3

**Table 1-2. Roll Compatibility**

Roll	Requires <sup>a</sup>	Optional <sup>b</sup>	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	


Notes:

- 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.
- 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.

# 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. Adding the Roll to a Running System

The Service Pack Roll can also be added to running frontend. Here's the procedure for an x86\_64 cluster (for an i386 cluster, substitute *x86\_64* with *i386*).

```
# cd /tmp
# wget ftp://ftp.rocksclusters.org/pub/rocks/rocks-5.0/service-pack-5.0.3-1.x86_64.disk1.iso
# rocks add roll service-pack*iso
# rocks enable roll service-pack
# cd /home/install
# rocks-dist dist
# kroll service-pack > /tmp/service-pack-install.sh
# sh /tmp/service-pack-install.sh
```

Then, reboot your frontend:

```
# init 6
```

When the frontend reboots, the Service Pack Roll is now installed. There is no need to reinstall your compute nodes.



# Chapter 3. Fixes Included in the Service Pack Roll

- Base Roll - An installation scalability fix. When multiple compute nodes are simultaneously reinstalled, there is a chance that some of the compute nodes will not complete their install. If you plug a monitor into a compute node in this state (or if you connect to the compute node's virtual console with `rocks-console`), you'll see a message indicating that the compute node can not download a specific package.

The fix is to aggressively retry when there is a package download failure.

- SGE Roll - A fix for the 'Job Queue' web page. The output of `'qstat -f'` changed between SGE 6.0 and 6.1u4 that caused no queued jobs to be reported on the 'Job Queue' ganglia web page.

The fix is to change the way the ganglia metric gets job queue info from SGE (the old method was `'qstat -f -xml'` and the new method is `'qstat -f -u \* -xml'`).

- Base Roll - A 411 scalability fix. In larger clusters, `'rocks sync users'` was not pushing all the user-related files (e.g., `/etc/passwd`, `/etc/shadow`, etc.) out to all the nodes. This was due to the fact that the 411 listener on each compute node, was rebroadcasting the 411 alert that it received from the frontend. On a vanilla Rocks cluster, there are 9 files that are pushed from the frontend to the cluster nodes when one executes `'rocks sync users'`. On a 100-node cluster without this fix that means that each compute node would receive 900 411 alert messages!

This fix removes the rebroadcasting of the 411 alert messages, so now each compute node only receives 9 alerts, regardless of the cluster size.

- Base Roll - Fix for Dell Virtual Floppy. On Dell machines that have Dell Remote Access Controllers (DRACs), the Dell Virtual Floppy would be recognized by anaconda as a hard disk. This would cause an exception that looks like:

```
Traceback (most recent call first):
  File "/tmp/product/autopart.py", line 1230, in doAutoPartition
    for f in initial_free[drive]:
  File "/usr/lib/anaconda/dispatch.py", line 201, in moveStep
    rc = stepFunc(self.anaconda)
  File "/usr/lib/anaconda/dispatch.py", line 124, in gotoNext
    self.moveStep()
  File "/usr/lib/anaconda/dispatch.py", line 223, in currentStep
    self.gotoNext()
  File "/usr/lib/anaconda/iw/release_notes.py", line 86, in resize
    (step, args) = self.anaconda.dispatch.currentStep()
  File "/usr/lib/anaconda/iw/release_notes.py", line 51, in __init__
    self.resize()
  File "/tmp/product/gui.py", line 1148, in __init__
    self.rnv = ReleaseNotesViewer(self.anaconda)
  File "/tmp/product/gui.py", line 977, in run
    self.icw = InstallControlWindow (self.anaconda)
  File "/usr/bin/anaconda", line 982, in ?
    anaconda.intf.run(anaconda)
KeyError: 'sdb'
```

The fix is to filter out hard disks that don't have 'media present' (which is how the Dell Virtual Floppy behaves).

# Appendix A. Rocks Copyright

Rocks(r)  
www.rocksclusters.org  
version 5.0 (V)

Copyright (c) 2000 - 2008 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.

