

Oracle® Exadata Storage Server Software

Release Notes

11g Release 1 (11.1.3)

E12062-03

May 2009

These release notes provide important information about Oracle Exadata Storage Server Software 11g Release 1 (11.1.3). Oracle Exadata Storage Server Software is called Exadata Cell or simply a cell.

After you review these release notes, read the *Oracle Exadata Storage Server Software User's Guide*.

See Also: Documentation updates for Oracle Exadata Storage Server Software, refer to Oracle*MetaLink* note 791275.1.

This release note contains the following topics:

- [Supported Environment and Preinstallation Requirements](#)
- [Installation and Configuration Issues](#)
- [Enhancements and Changes included in Oracle Exadata Storage Server Software 11g Release 1 \(11.1.3\) Patch Set 2 \(11.1.3.2\)](#)
- [Known Issues](#)
- [Documentation Accessibility](#)

1 Supported Environment and Preinstallation Requirements

This section describes the supported environment and preinstallation requirements.

- [Database and ASM Instances](#)
- [Exadata Cell](#)

1.1 Database and ASM Instances

This section describes preinstallation requirements for the database and ASM instance.

- Use the Oracle Database 11g Release 1 (11.1.0.7) Patch Set 1 with Oracle Exadata Storage Server Software code.

1.2 Exadata Cell

The preconfigured systems that have been optimized for use. The storage cell has been imaged with the necessary software. Do not change or customize the storage cell. If you do so, then you risk losing your changes and possibly voiding your support contract.

2 Installation and Configuration Issues

This section describes the installation and configuration issues.

- [ASM Instance](#)
- [Database Instance](#)
- [Oracle Real Application Clusters](#)
- [Cluster Ready Services](#)

For information about installation and configuration, see *Oracle Exadata Storage Server Software User's Guide*.

2.1 ASM Instance

This section describes installation and configuration issues for the ASM instance.

- When installing the Oracle database, the Automatic Storage Management (ASM) storage option must be selected.
- Turning on the Linux firewall (`iptables`) and configuring it incorrectly can cause the ASM instance to fail to recognize disks from Cell Server (CELLSRV), or in some situations, the discovery process halts. The firewall should be turned off on the database machine and storage cell.

Use the following commands in a Linux environment to:

- To check the firewall status, use the following command:

```
/etc/init.d/iptables status
```

- To stop the firewall, use the following command:

```
/etc/init.d/iptables stop
```

- To disable the firewall at startup, use the following command:

```
/sbin/chkconfig iptables off
```

2.2 Database Instance

This section describes installation and configuration issues for the database instance.

- For this release, the setting for the `COMPATIBLE` database initialization parameter is `'11.1.0.7'`.

During Oracle Database installation, Database Configuration Assistant (DBCA) 11.1.0.7 does not default to `COMPATIBLE` equal to 11.1.0.7, so you must set the value on the Initialization Parameters page.

2.3 Oracle Real Application Clusters

This section describes installation and configuration issues for Oracle Real Application Clusters (Oracle RAC).

- For an Oracle RAC cluster, you must shut down the entire cluster and then restart it after the database and cell software have been installed and the cell configuration files are in place.

2.4 Cluster Ready Services

This section describes installation and configuration issues for Oracle Clusterware Cluster Ready Services (CRS).

- When installing Oracle Clusterware software, ensure the private IP address is specified as the same as the InfiniBand IP address used by the storage cells to send data to the database server host. The IP address is defined in the `cellinit.ora` file.

3 Enhancements and Changes included in Oracle Exadata Storage Server Software 11g Release 1 (11.1.3) Patch Set 2 (11.1.3.2)

Oracle Exadata Storage Server Software 11g Release 1 (11.1.3) Patch Set 2 (11.1.3.2) set includes enhancements and changes to Oracle Exadata Storage Server Software. This section contains the following topics:

- [Larger Maximum File Size and Improved Performance for Local File System on Database Node](#)
- [System Configuration for Network and Related Properties](#)
- [Updated Firmware for 450GB SAS Drives](#)

3.1 Larger Maximum File Size and Improved Performance for Local File System on Database Node

The performance and responsiveness of the local file system is substantially improved when working with very large files on local file system (/) due to the new 4KB file system block size on database nodes. This change addresses the Oracle Real Application Clusters (Oracle RAC) node evictions issue that had been noticed while doing sustained large writes through very large files on local file system.

This change also increases the maximum file size for local file system from 16GB to 2TB.

3.2 System Configuration for Network and Related Properties

Restrictions on manual configuration changes to database nodes, such as hostname or IP address, are removed. With 11g Release 1 (11.1.3), only the first-time configuration is carried out on the database node.

Use your own methods or tools to do any changes after the first-time configuration to the node, such as to IP address, InfiniBand (IB) IP address, hostname, time zone, NTP servers, and DNS Name servers.

See Also: There have been simplifications for changing IP addresses and IB addresses on Exadata Cells. Refer to chapter 4 of the *Oracle Exadata Storage Server Software User's Guide* for details.

3.3 Updated Firmware for 450GB SAS Drives

Exadata Cell BIOS, disk controller and the disk firmware for 450GB SAS drive-based cells is updated to address hang, false failures of disks, and other performance issues.

4 Known Issues

This section contains other known issues.

- [Exadata Cell Issues](#)
- [CellCLI Command Line Tool Issues](#)
- [ASM Instance](#)
- [Database Instance](#)
- [Cluster Ready Services](#)
- [Ignorable Error Messages](#)

4.1 Exadata Cell Issues

This section describes general Exadata Cell issues.

- To utilize bloom predicate pushdown, the tables on both sides of a `JOIN` should be created with the `PARALLEL` clause.

4.2 CellCLI Command Line Tool Issues

This section describes CellCLI issues.

- All attributes for an Exadata Cell object can be displayed with the `DESCRIBE` command. However, some of the attribute values, such as the value for the `errorCount` attribute, are not updated.

4.3 ASM Instance

This section describes general ASM instance issues.

- `ALTER DISKGROUP ... ONLINE ... SQL` statement must use the `WAIT` option.
- `DROP DISKGROUP` fails if there is a staleness registry. The workaround is to use `DROP DISKGROUP...FORCE`. Note that a staleness registry is created when there is an offline disk in the disk group.

4.4 Database Instance

There are no general database instance issues.

4.5 Cluster Ready Services

There are no general Oracle Clusterware Cluster Ready Services (CRS) issues.

4.6 Ignorable Error Messages

This section lists several messages that may appear on the console, in the `/var/log/messages` directory, or other log files. These are messages that are not known to have any unfavorable effects.

This section contains the following topics:

- [Ignorable Error Messages During Reimaging and Patching](#)
- [Ignorable Error Messages on the Console Related to InfiniBand](#)
- [Ignorable Error Messages on Database Nodes](#)
- [Ignorable Error Messages in Log Files and on Console](#)
- [Ignorable Error Messages when Applying Patch to Oracle Database and Exadata Cell](#)
- [Ignorable Error Messages on Database and Cell Nodes when Applying Patch Set](#)
- [Ignorable Error Messages on Database and Cell Nodes when Applying Patch Set or Reimaging](#)
- [Ignorable Error Messages while Booting or Shutting Down](#)
- [Ignorable Error Messages During Patching to 11g Release 1 \(11.1.3\)](#)

4.6.1 Ignorable Error Messages During Reimaging and Patching

The following error messages may appear during reimaging and patching, and appear in the `/var/log/messages` directory logs. These messages can be ignored.

```
"cciss/cXdYpZ Invalid partition table"
    The above mescell will repeat several times with X,Y, and Z are some
integers
    - X is the P400 smart array disk controller slot number
    - Y is the disk number starting with 0 for the slot
    - Z is the partition number on the disk
```

```
"RAID1 conf printout:..."
```

The preceding error message appears in the software RAID printout.

4.6.2 Ignorable Error Messages on the Console Related to InfiniBand

The following error messages may appear on the console or in the `/var/log/messages` logs. These messages can be ignored.

```
ib0: multicast join failed for ff12:401b:ffff:0000:0000:0000:ffff:ffff,
status -11
ib1: multicast join failed for ff12:401b:ffff:0000:0000:0000:ffff:ffff,
status -11
```

```
kernel: bonding: bond0: Warning: enslaved VLAN challenged slave ib0. Adding
VLANs
will be blocked as long as ib0 is part of bond bond0
kernel: bonding: bond0: Warning: The first slave device you specified does
not
support setting the MAC address. This bond MAC address would be that of the
active
slave.
```

```

kernel: bonding: bond0: Warning: enslaved VLAN challenged slave ib1. Adding
VLANs
will be blocked as long as ib1 is part of bond bond0

kernel: ib0: multicast join failed for
ff12:401b:ffff:0000:0000:0000:ffff:ffff,
status -11

kernel: bonding: Warning: either miimon or arp_interval and arp_ip_target
module
parameters must be specified, otherwise bonding will not detect link
failures! see
bonding.txt for details.

kernel: bonding: bond0: Warning: The first slave device you specified does
not
support setting the MAC address. This bond MAC address would be that of the
active
slave.

snmpd[4038]: error finding row index in _ifXTable_container_row_restore <Time
adjusted by NTPD> so these messages are actually later in boot process

kernel: bonding: bond0: Warning: the permanent HWaddr of ib0 -
80:00:00:48:FE:80 -
is still in use by bond0. Set the HWaddr of ib0 to a different address to
avoid
conflicts.

kernel: bonding: bond0: Warning: enslaved VLAN challenged slave ib1. Adding
VLANs
will be blocked as long as ib1 is part of bond bond0

snmpd[6136]: netsnmp_assert index == tmp failed if-mib/data_
access/interface.c:
467_access_interface_entry_save_name()
snmpd[4108]: netsnmp_assert index == tmp failed if-mib/data_
access/interface.c:
467_access_interface_entry_save_name()
snmpd[4105]: error finding row index in _ifXTable_container_row_restore
snmpd[4103]: netsnmp_assert index == tmp failed if-mib/data_
access/interface.c:
467_access_interface_entry_save_name()
snmpd[4100]: error finding row index in _ifXTable_container_row_restore

snmpd[6136]: netsnmp_assert index == tmp failed if-mib/data_
access/interface.c:467_access_interface_entry_save_name()
snmpd[4038]: error finding row index in _ifXTable_container_row_restore
<Time adjusted by NTPD> so these messages are actually later in boot process

```

4.6.3 Ignorable Error Messages on Database Nodes

The following error messages may appear for the database nodes. These messages can be ignored.

```

kernel: ACPI Error (psparse-0537): Method parse/execution failed [\_PR_
.CPU7._CST]
(Node ffff81082fd712f0), AE_NOT_FOUND
kernel: cciss: MSI init failed

```

4.6.4 Ignorable Error Messages in Log Files and on Console

The following error messages may appear in the `/var/log/messages` directory or on the console when patching the system:

```
ipmi_si(SI_CHECK_BMC): Failed to get Global Enables 0xc6.
OsKcsExecCmd: IPMI NetFN 0x36 CMD: 0x2 has timed out!
OsKcsExecCmd: IPMI NetFN 0x36 CMD: 0x2 has timed out!
OsKcsExecCmd: IPMI NetFN 0x36 CMD: 0x2 has timed out!
OsKcsExecCmd: IPMI NetFN 0x36 CMD: 0x2 has timed out!
warning: many lost ticks.
Your time source seems to be instable or some driver is hogging interrupts
rip s_show+0x167/0x293
```

```
process 'sysctl' is using deprecated sysctl (syscall)
net.ipv6.neigh.eth0.retrans_
time;
Use net.ipv6.neigh.eth0.retrans_time_ms instead.
process 'cp' is using deprecated sysctl (syscall) net.ipv6.neigh.eth0.base_
reachable_time;
Use net.ipv6.neigh.eth0.base_reachable_time_ms instead.
```

```
EXT3-fs warning: maximal mount count reached, running e2fsck is recommended
EXT3-fs warning: maximal mount count reached, running e2fsck is recommended
```

4.6.5 Ignorable Error Messages when Applying Patch to Oracle Database and Exadata Cell

The following error messages may appear when applying the patch set to Oracle Database and Exadata Cells. These messages can be ignored.

```
Patching files in / ...
Patching /boot
Extracting ofed/rds files in /install/pre/ofed ...
bonding: bond0: doing slave updates when interface is down.
bonding: bond0: Warning: the permanent HWaddr of ib0 - 80:00:00:48:FE:80 - is
still in use by bond0. Set the HWaddr of ib0 to a different address to avoid
conflicts.
bonding: bond0: doing slave updates when interface is down.
dev_mc_discard: multicast leakage! dmi_users=1
dev_mc_discard: multicast leakage! dmi_users=1
dev_mc_discard: multicast leakage! dmi_users=1
Installing rds/ofed packages ...
```

4.6.6 Ignorable Error Messages on Database and Cell Nodes when Applying Patch Set

The following error messages may appear in the `/var/log/messages` directory or on the console while applying the patch set. These messages can be ignored.

```
Memory for crash kernel (0x0 to 0x0) not within permissible range

insmod cannot insert '/lib/modules/2.6.18-53.1.21.2.1.e15/
kernel/drivers/usb/host/uhci-hcd.ko':
File exists (-1) File exists
```

4.6.7 Ignorable Error Messages on Database and Cell Nodes when Applying Patch Set or Reimaging

The following error messages may appear in the `/var/log/messages` directory or on the console while applying the patch set or reimaging. These messages can be ignored.

```
ln: creating symbolic link '/opt/hp/hpsmh/data/htdocs/ADU' to
'/opt/hp/hpadu/HTML': No such file or directory
```

```
chown: 'hpsmh:hpsmh': invalid user
```

```
ln: creating symbolic link '/opt/hp/hpsmh/webapp/ADU.xml.inactive' to
'/opt/hp/hpadu/bld/adu.xml': No such file or directory
```

4.6.8 Ignorable Error Messages while Booting or Shutting Down

The following error messages may appear in the `/var/log/messages` directory or on the console while booting or shutting down. These messages can be ignored:

```
Setting up InfiniBand network interfaces:
/etc/sysconfig/network-scripts/ifup-eth:
line 109: /sys/class/net/bond0/bonding/
slaves: No such file or directory
```

```
Bringing up interface ib0: /etc/sysconfig/network-scripts/ifup-eth: line 109:
/sys/class/net/bond0/bonding/slaves: No such file or directory
```

```
Shutting down interface bond0: bonding: bond0: Removing slave ib0
bonding: bond0: Warning: the permanet HWaddr of ib0 - 80:00:00:48:FE:80 -
is still in use by bond0.
```

```
Set HWaddr of ib0 to a different address to avoid conflicts.
```

```
Starting killall: Stopping ipmi drivers: may be in use
```

```
The poll API returned unexpected erro: Error Code: 29 - Interrupted system
call
```

```
Loading ata_piix.ko module
usb 4-1: device not accepting address 2, error -71
```

4.6.9 Ignorable Error Messages During Patching to 11g Release 1 (11.1.3)

The following error messages may appear in the `/var/log/messages` directory or on the console while applying the patch set. These messages can be ignored:

```
Starting openibd: Loading Mellanox MLX4 HCA driver [FAILED]
Loading HCA driver and Access Layer: [FAILED]
Please open an issue in the http://openib.org/bugzilla and attach
/tmp/ib_debug_info.log
```

```
bonding: unable to remove non-existent slave ib0 for bond bond0.
bonding: unable to remove non-existent slave ib1 for bond bond0.
bonding: bond0: Warning: enslaved VLAN challenged slave ib0. Adding VLANs wi
e blocked as long as ib0 is part of bond bond0
bonding: bond0: Error: The slave device you specified does not support
setting the MAC address. Your kernel
likely does not support slave devices.
bonding: bond0: Warning: enslaved VLAN challenged slave ib1. Adding VLANs wi
e blocked as long as ib1 is part of bond bond0
```


bonding: bond0: Error: The slaved device you specified does not support setting the MAC address.
Your kernel likely does not support slave devices.

tar: <some file name>: time stamp <some time is XXXX in the future

5 Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at <http://www.fcc.gov/cgb/consumerfacts/trs.html>, and a list of phone numbers is available at <http://www.fcc.gov/cgb/dro/trsphonebk.html>.

Oracle Exadata Storage Server Release Software Notes, 11g Release 1 (11.1.3)
E12062-03

Copyright © 2008, 2009, Oracle. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such

applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners. The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.