

**Oracle® VM Server for SPARC 3.1.1.2,  
3.1.1.1, 3.1.1, and 3.1 Release Notes**

Copyright © 2007, 2014, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

---

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

# Contents

---

<b>Preface</b> .....	5
<b>1 Oracle VM Server for SPARC 3.1.1.2, 3.1.1.1, 3.1.1, and 3.1 Release Notes</b> .....	7
Oracle VM Server for SPARC 3.1.1.2 Maintenance Update .....	8
Oracle VM Server for SPARC 3.1.1.1 Maintenance Update .....	8
What's New in This Release .....	8
What's New in the Oracle VM Server for SPARC 3.1.1.1 Maintenance Update .....	9
What's New in the Oracle VM Server for SPARC 3.1.1 Release .....	9
What's New in the Oracle VM Server for SPARC 3.1 Release .....	10
System Requirements .....	11
Supported Platforms .....	11
Required Software and Patches .....	13
Related Software .....	25
Software That Can Be Used With the Oracle VM Server for SPARC Software .....	25
System Controller Software That Interacts With Oracle VM Server for SPARC .....	26
Optional Software .....	27
Upgrading to the Current Oracle VM Server for SPARC Software .....	27
Upgrading to the Oracle VM Server for SPARC 3.1.1.1 Software .....	27
Upgrading to the Oracle VM Server for SPARC 3.1.1 Software .....	28
Upgrading to the Oracle VM Server for SPARC 3.1 Software .....	29
Deprecated Oracle VM Server for SPARC Features .....	29
Known Issues .....	30
General Issues .....	30
Domain Migration Restrictions .....	42
Oracle VM Server for SPARC MIB Issues .....	45
SR-IOV Issues .....	45
Bugs Affecting the Oracle VM Server for SPARC Software .....	49
Documentation Issues .....	99

Resolved Issues .....	100
Resolved Issues in the Oracle VM Server for SPARC 3.1.1.2 Release .....	100
Resolved Issues in the Oracle VM Server for SPARC 3.1.1.1 Release .....	100
Resolved Issues in the Oracle VM Server for SPARC 3.1.1 Release .....	101
Resolved Issues in the Oracle VM Server for SPARC 3.1.0.1 Release .....	102
Resolved Issues in the Oracle VM Server for SPARC 3.1 Release .....	102

# Preface

---

- **Overview** – Includes information about this release of the Oracle VM Server for SPARC software, such as changes for this release, supported platforms, a matrix of required software and patches, and bugs affecting this software.
- **Audience** – System administrators who manage virtualization on SPARC servers
- **Required knowledge** – System administrators on these servers must have a working knowledge of UNIX systems and the Oracle Solaris operating system (Oracle Solaris OS)

## Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at [http://www.oracle.com/pls/topic/lookup?ctx=product\\_intuitive\\_ID](http://www.oracle.com/pls/topic/lookup?ctx=product_intuitive_ID).

## Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

## Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.



# Oracle VM Server for SPARC 3.1.1.2, 3.1.1.1, 3.1.1, and 3.1 Release Notes

---

These release notes contain changes for this release, a list of supported platforms, a matrix of required software and patches, and other pertinent information, including bugs that affect the Oracle VM Server for SPARC 3.1.1.2, Oracle VM Server for SPARC 3.1.1.1, Oracle VM Server for SPARC 3.1.1, and Oracle VM Server for SPARC 3.1 software.

---

**Note** – Oracle VM Server for SPARC features are added and maintained on the supported hardware platforms listed in [“Supported Platforms” on page 11](#). However, new features will not be added and existing features will not be maintained on hardware platforms that have been removed from the list.

As a rule, new Oracle VM Server for SPARC features and functionality are made available for all price-listed, supported T-Series and M-Series servers from Oracle and Fujitsu M10 systems at the time that the Oracle VM Server for SPARC software is released and not for SPARC based systems that have already passed their last-order date.

---

---

**Note** – The Oracle VM Server for SPARC 3.1.1.2 software includes all of the bug fixes since the Oracle VM Server for SPARC 3.0 release and is the patch update for the existing releases. If you are running the earlier Oracle VM Server for SPARC 3.0 or Oracle VM Server for SPARC 3.1 software, upgrade to the latest Oracle VM Server for SPARC 3.1.1.2 update to obtain additional bug fixes.

---

## Oracle VM Server for SPARC 3.1.1.2 Maintenance Update

The following sections cover the Oracle VM Server for SPARC 3.1.1.2 maintenance release:

- “Bugs Affecting the Oracle VM Server for SPARC 3.1.1.2 Software” on page 49
- “Resolved Issues in the Oracle VM Server for SPARC 3.1.1.2 Release” on page 100

The following sections include information that also applies to the Oracle VM Server for SPARC 3.1.1.2 software but has not changed since the Oracle VM Server for SPARC 3.1.1.1 maintenance release:

- “What's New in the Oracle VM Server for SPARC 3.1.1.1 Maintenance Update” on page 9
- “Required Oracle Solaris OS Versions for the Oracle VM Server for SPARC 3.1.1.1 Maintenance Update” on page 13
- “Required Software to Enable the Latest Oracle VM Server for SPARC Features” on page 16
- “Required System Firmware Patches” on page 17
- “Minimum Version of Software Required” on page 17
- “Upgrading to the Oracle VM Server for SPARC 3.1.1.1 Software” on page 27

## Oracle VM Server for SPARC 3.1.1.1 Maintenance Update

This section includes links to sections in this book that cover the Oracle VM Server for SPARC 3.1.1.1 maintenance release:

- “What's New in the Oracle VM Server for SPARC 3.1.1.1 Maintenance Update” on page 9
- “Required Oracle Solaris OS Versions for the Oracle VM Server for SPARC 3.1.1.1 Maintenance Update” on page 13
- “Required Software to Enable the Latest Oracle VM Server for SPARC Features” on page 16
- “Required System Firmware Patches” on page 17
- “Minimum Version of Software Required” on page 17
- “Upgrading to the Oracle VM Server for SPARC 3.1.1.1 Software” on page 27
- “Bugs Affecting the Oracle VM Server for SPARC 3.1.1.1 Software” on page 50
- “Resolved Issues in the Oracle VM Server for SPARC 3.1.1.1 Release” on page 100

## What's New in This Release

For information about the features introduced in all versions of the Oracle VM Server for SPARC (Logical Domains) software, see [What's New in Oracle VM Server for SPARC Software](http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html) (<http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html>).

## What's New in the Oracle VM Server for SPARC 3.1.1.1 Maintenance Update

The changes for this 3.1.1.1 maintenance update of the Oracle VM Server for SPARC software are as follows. Note that each feature is available on all supported platforms unless otherwise stated.

- Provides support for private virtual LANs. See “Using Private VLANs” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

This feature requires at least the Oracle Solaris 11.2.4.0.0 (SRU 4) OS.

- Increases the number of LDC endpoints per guest domain on SPARC T4, SPARC T5, SPARC M5, and SPARC M6 systems. See “Using Logical Domain Channels” in *Oracle VM Server for SPARC 3.1 Administration Guide*.
- Adds support for dynamically assigning PCIe endpoint devices to Fujitsu M10 systems. See the *Fujitsu M10 Systems Product Notes*.

## What's New in the Oracle VM Server for SPARC 3.1.1 Release

The major changes for this release of the Oracle VM Server for SPARC 3.1.1 software are as follows. Note that each feature is available on all supported platforms unless otherwise stated.

- Adds support for Fibre Channel SR-IOV. See “Creating an I/O Domain by Assigning PCIe SR-IOV Virtual Functions” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

This feature is supported only on SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

- Adds network bandwidth controls. See “Controlling the Amount of Physical Network Bandwidth That Is Consumed by a Virtual Network Device” in *Oracle VM Server for SPARC 3.1 Administration Guide*.
- Adds ability to create or destroy multiple virtual functions simultaneously. See the `ldm(1M)` man page.

# What's New in the Oracle VM Server for SPARC 3.1 Release

The major changes for this release of the Oracle VM Server for SPARC 3.1 software are as follows. Note that each feature is available on all supported platforms unless otherwise stated.

- Enhances dynamic I/O virtualization for SR-IOV. See [Chapter 6, “Setting Up I/O Domains,” in \*Oracle VM Server for SPARC 3.1 Administration Guide\*](#) and [“Fujitsu M10 System Has Different SR-IOV Feature Limitations” on page 48](#).

This feature is supported only on SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

- Adds support for non-primary domains to act as root domains. See [Chapter 6, “Setting Up I/O Domains,” in \*Oracle VM Server for SPARC 3.1 Administration Guide\*](#).

This feature is supported only on SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms, and Fujitsu M10 systems.

- Enhances the dynamic reconfiguration operation to act on all types of root domains. See [Chapter 6, “Setting Up I/O Domains,” in \*Oracle VM Server for SPARC 3.1 Administration Guide\*](#).

This feature is supported only on SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

- Adds the Fault Management Architecture (FMA) blacklisting feature. When FMA detects faulty CPU or memory resources, Oracle VM Server for SPARC places them on a blacklist. A faulty resource that is on the blacklist cannot be reassigned to any domains until FMA marks it as being repaired. See [Chapter 12, “Handling Hardware Errors,” in \*Oracle VM Server for SPARC 3.1 Administration Guide\*](#).

This feature is supported only on SPARC T5, SPARC M5, and SPARC M6 platforms.

- Adds a recovery mode feature that automatically recovers domain configurations that cannot be booted because of faulty or missing resources. See [Chapter 12, “Handling Hardware Errors,” in \*Oracle VM Server for SPARC 3.1 Administration Guide\*](#).

This feature is supported only on SPARC T5, SPARC M5, and SPARC M6 platforms.

- Adds cross-CPU migration support for the SPARC T5, SPARC M5, and SPARC M6 platforms. This feature enables you to perform a migration even if the CPU type of the source and target systems is not identical. Use the `ldm set -domain` command to set the `cpu - arch` property.

- Adds InfiniBand device support to the SR-IOV feature. See [Chapter 6, “Setting Up I/O Domains,” in \*Oracle VM Server for SPARC 3.1 Administration Guide\*](#).

This feature is supported only on SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

- Adds the `ldmpower` command to show per-domain, power-consumption information. See “Viewing Power-Consumption Data” in *Oracle VM Server for SPARC 3.1 Administration Guide* and the `ldmpower(1M)` man page.  
This feature is supported only on SPARC T3, SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms.
- Adds support for VNICs on virtual networks. See the `ldm(1M)` man page.
- Significant networking performance improvements. See Chapter 8, “Using Virtual Networks,” in *Oracle VM Server for SPARC 3.1 Administration Guide*.
- Bug fixes.

## System Requirements

This section contains system requirements for running the Oracle VM Server for SPARC software.

### Supported Platforms

When more than one version of the Oracle VM Server for SPARC software is supported on a hardware platform, bug fixes apply only to the latest version of the software. To receive Premier Support, you must use the latest Oracle VM Server for SPARC software.

You can find the platform documentation on the [Oracle Technology Network](http://www.oracle.com/technetwork/documentation/sparc-tseries-servers-252697.html) (<http://www.oracle.com/technetwork/documentation/sparc-tseries-servers-252697.html>). You can also find information about the software stacks for the various platforms on the [Sun System Software Stacks page](http://www.oracle.com/technetwork/systems/software-stacks/stacks/index.html) (<http://www.oracle.com/technetwork/systems/software-stacks/stacks/index.html>).

The Oracle VM Server for SPARC 3.1 software is supported on the following platforms:

- **Fujitsu M-Series servers**
  - Fujitsu M10 system (refer to *Fujitsu M10 Systems Product Notes*)  
For more information about features that are specific to the Fujitsu M10 systems, see *Fujitsu M10 Systems System Operation and Administration Guide* in the Product Notes for your model at <http://www.fujitsu.com/global/services/computing/server/sparc/downloads/manual/>.
- **SPARC M-Series servers**
  - SPARC M6-32 server (refer to *SPARC M5-32 and SPARC M6-32 Servers Product Notes*)
  - SPARC M5-32 server (refer to *SPARC M5-32 Server Product Notes*)
- **SPARC T5 servers**

- SPARC T5-1B server (refer to *SPARC T5-1B Server Product Notes*)
- SPARC T5-2 server (refer to *SPARC T5-2 Server Product Notes*)
- SPARC T5-4 server (refer to *SPARC T5-4 Server Product Notes*)
- SPARC T5-8 server (refer to *SPARC T5-8 Server Product Notes*)
- **SPARC T4 servers**
  - SPARC T4-1 server (refer to *SPARC T4-1 Server Product Notes*)
  - SPARC T4-2 server (refer to *SPARC T4-2 Server Product Notes*)
  - SPARC T4-4 server (refer to *SPARC T4-4 Server Product Notes*)
  - SPARC T4-1B server (refer to *SPARC T4-1B Server Product Notes*)
  - Netra SPARC T4-1 server (refer to *Netra SPARC T4-1 Server Product Notes*)
  - Netra SPARC T4-2 server (refer to *Netra SPARC T4-2 Server Product Notes*)
  - Netra SPARC T4-1B server (refer to *Netra SPARC T4-1B Server Product Notes*)
- **SPARC T3 servers**
  - SPARC T3-1 server (refer to *SPARC T3-1 Server Product Notes*)
  - SPARC T3-2 server (refer to *SPARC T3-2 Server Product Notes*)
  - SPARC T3-4 server (refer to *SPARC T3-4 Server Product Notes*)
  - SPARC T3-1B server (refer to *SPARC T3-1B Server Module Product Notes*)
  - Netra SPARC T3-1 server (refer to *Netra SPARC T3-1 Server Product Notes*)
  - Netra SPARC T3-1B server (refer to *Netra SPARC T3-1B Server Product Notes*)
  - Netra SPARC T3-1BA server (refer to *Netra SPARC T3-1BA Server Product Notes*)
- **UltraSPARC T2 Plus servers**
  - Oracle's Sun SPARC Enterprise T5140 and T5240 servers (refer to *Sun SPARC Enterprise T5140 and T5240 Servers Administration Guide*)
  - Oracle's Sun SPARC Enterprise T5440 server (refer to *Sun SPARC Enterprise T5440 Server Administration Guide*)
  - Oracle's Sun Blade T6340 server module (refer to *Sun Blade T6340 Server Module Product Notes*)
  - Oracle's Netra T5440 server (refer to *Sun Netra T5440 Server Product Notes*)
  - Oracle's Sun Netra T6340 Server Module (refer to *Sun Netra T6340 Server Module Product Notes*)
- **UltraSPARC T2 servers**
  - Oracle's Sun SPARC Enterprise T5120 and T5220 servers (refer to *Sun SPARC Enterprise T5120 and T5220 Servers Administration Guide*)
  - Oracle's Sun Blade T6320 server module (refer to *Sun Blade T6320 Server Module Product Notes*)
  - Oracle's Netra T5220 server (refer to *Sun Netra T5220 Server Product Notes*)
  - Oracle's Netra CP3260 Blade (refer to *Netra CP3260 Blade Server Product Notes*)

---

## Required Software and Patches

This section lists the required software and patches for use with the current Oracle VM Server for SPARC software.



---

**Caution** – Do not downgrade to older versions of individual software and firmware components. Such downgrades are not recommended and might lead to unexpected behavior and failures.

---

### Required Oracle Solaris OS Versions

#### Required Oracle Solaris OS Versions for the Oracle VM Server for SPARC 3.1.1.1 Maintenance Update

To use all features of the Oracle VM Server for SPARC 3.1.1.1 maintenance update, the operating system (OS) on all domains should be at least the Oracle Solaris 10 1/13 OS with Oracle VM Server for SPARC 3.1.1.1 patch ID 150817-03 and the required patches shown in [Table 1-1](#) or the Oracle Solaris 11.2.2.0.5 OS. This OS can be either an initial installation of or an upgrade to the appropriate OS.

The PVLAN feature requires at least the Oracle Solaris 11.2.4.0.0 (SRU 4) OS running on the service domain. This OS is due for release in November 2014.

---

**Note** – On Oracle Solaris 10, you must disable the `ldmd` daemon before you apply Oracle VM Server for SPARC 3.1.1.1 patch 150817-03. Re-enable the daemon after the patch is applied. See [“Upgrading to the Oracle VM Server for SPARC 3.1.1.1 Software” on page 27](#).

---

The following table shows the patches that must be applied to the Oracle Solaris 10 1/13 OS to include the Oracle VM Server for SPARC 3.1.1.1 functionality. For information about the minimum Oracle Solaris OS versions for supported server platforms, see the data sheet for your server platform at <http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html>.

---

**Note** – On an Oracle Solaris 10 system, you can upgrade to the Oracle VM Server for SPARC 3.1.1.1 maintenance update directly from either the Oracle VM Server for SPARC 3.1.1 or the Oracle VM Server for SPARC 3.1 versions by applying the Oracle VM Server for SPARC 3.1.1.1 patch.

---

**TABLE 1-1** Oracle VM Server for SPARC 3.1.1.1: Patches for Older Oracle Solaris OS Versions and Domains Requiring a Patch

Patch ID	Control Domain	Service Domain	I/O Domain	Guest Domain
125555-15 (SunOS 5.10: patch behavior patch)	X	X	X	
146582-03 (SunOS 5.10: fmd patch)	X	X	X	
148322-08 (SunOS 5.10: ixgbe patch)	X	X	X	
148324-07 (SunOS 5.10: ixgbev f patch)	X	X	X	
148888-05 (SunOS 5.10: kernel patch)	X	X	X	
149173-04 (SunOS 5.10: emlxs driver patch)	X	X	X	X
150107-01 (SunOS 5.10: ds patch)	X	X	X	X
150400-11 (SunOS 5.10: kernel patch)	X	X	X	X
150435-02 (SunOS 5.10: placeholder patch)	X	X	X	X
150840-02 (SunOS 5.10: platsvc patch)	X	X	X	X

**Note** – This patch list includes the minimum required patch revisions. You can install later revisions of the same patch.

### Required Oracle Solaris OS Versions for Oracle VM Server for SPARC 3.1.1

To use all features of the Oracle VM Server for SPARC 3.1.1 software, the operating system (OS) on all domains should be at least the Oracle Solaris 11.1.17.4.0 OS or the Oracle Solaris 10 1/13 OS plus the Oracle VM Server for SPARC 3.1.1 patch ID 150817-02 and the required patches in [Table 1-2](#). This OS can be either an initial installation of or an upgrade to the appropriate OS.

**Note** – You must disable the `ldmd` daemon before you apply Oracle VM Server for SPARC 3.1.1 patch 150817-02. Re-enable the daemon after the patch is applied. See [“Upgrading to the Oracle VM Server for SPARC 3.1.1 Software” on page 28](#).

The following table shows the patches that must be applied to the Oracle Solaris 10 1/13 OS to include the Oracle VM Server for SPARC 3.1.1 functionality. For information about the minimum Oracle Solaris OS versions for supported server platforms, see the data sheet for your server platform at <http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html>.

**TABLE 1-2** Oracle VM Server for SPARC 3.1.1: Patches for Older Oracle Solaris OS Versions and Domains Requiring a Patch

Patch ID	Control Domain	Service Domain	I/O Domain	Guest Domain
125555-15 (SunOS 5.10: patch behavior patch)	X	X	X	
146582-03 (SunOS 5.10: fmd patch)	X	X	X	
148322-08 (SunOS 5.10: ixgbe patch)	X	X	X	
148324-07 (SunOS 5.10: ixgbev f patch)	X	X	X	
148888-05 (SunOS 5.10: kernel patch)	X	X	X	
149173-04 (SunOS 5.10: emlxs driver patch)	X	X	X	X
150107-01 (SunOS 5.10: ds patch)	X	X	X	X
150400-11 (SunOS 5.10: kernel patch)	X	X	X	X
150435-02 (SunOS 5.10: placeholder patch)	X	X	X	X
150840-02 (SunOS 5.10: platsvc patch)	X	X	X	X

**Note** – This patch list includes the minimum required patch revisions. You can install later revisions of the same patch.

## Required Oracle Solaris OS Versions for Oracle VM Server for SPARC 3.1

To use all features of the Oracle VM Server for SPARC 3.1 software, the operating system (OS) on all domains should be at least the Oracle Solaris 10 1/13 OS or the Oracle Solaris 11.1.10.5.0 OS. This OS can be either an initial installation of or an upgrade to the Oracle Solaris 10 1/13 OS or the Oracle Solaris 11.1.10.5.0 OS.

The following table shows the patches that must be applied to the Oracle Solaris 10 1/13 OS to include the Oracle VM Server for SPARC 3.1 functionality. For information about the minimum Oracle Solaris OS versions for supported server platforms, see the data sheet for your server platform at <http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html>.

**TABLE 1-3** Oracle VM Server for SPARC 3.1: Patches for Older Oracle Solaris OS Versions and Domains Requiring a Patch

Patch ID	Control Domain	Service Domain	I/O Domain	Guest Domain
146582-03 (SunOS 5.10: fmadm patch)	X	X	X	
148322-08 (SunOS 5.10: ixgbe patch)	X	X	X	
148324-07 (SunOS 5.10: ixgbev f patch)	X	X	X	
148888-05 (SunOS 5.10: kernel patch)	X	X	X	
150107-01 (SunOS 5.10: ds patch)	X	X	X	X
150400-02 (SunOS 5.10: kernel patch)	X	X	X	X
150840-02 (SunOS 5.10: platsvc patch)	X	X	X	X

**Note** – This patch list includes the minimum required patch revisions. You can install later revisions of the same patch.

## Required Software to Enable the Latest Oracle VM Server for SPARC Features

To enable all the Oracle VM Server for SPARC 3.1.1.1, Oracle VM Server for SPARC 3.1.1, or Oracle VM Server for SPARC 3.1 features, you must run the required system firmware versions on the platforms listed in the following table.

**TABLE 1-4** Oracle VM Server for SPARC 3.1.1 and 3.1: Required System Firmware Versions

Platform Type	3.1.1.1	3.1.1	3.1
Fujitsu M10 System	XCP2210	XCP2210	XCP2051
SPARC M6 Server	9.2.1.c <sup>1</sup>	9.1.2.d	9.1.0.g
SPARC M5 Server	9.2.1.c <sup>1</sup>	9.1.2.d	9.1.0.f
SPARC T5 Server	9.2.1.b <sup>1</sup>	9.1.2.d	9.1.0.b
SPARC T4 Server	8.5.1.b <sup>1</sup>	8.4.2.c	8.4.0.a
SPARC T3 Server	8.3	8.3	8.3
SPARC T2 Plus Server	7.4.5	7.4.5	7.4.5
SPARC T2 Server	7.4.5	7.4.5	7.4.5

<sup>1</sup> Only update to this system firmware version if you want to increase the number of LDC endpoints on guest domains. Using this firmware might result in live migration problems, see [“Live Migration Might Fail With Unable to restore ldc resource state on target Domain Migration of LDom failed”](#) on page 50.

## Required System Firmware Patches

This table lists the specific patch numbers that deliver the required versions that are listed in [Table 1–4](#).



**Caution** – Only install to the 3.1.1.1 system firmware patch listed in this table if you want to increase the number of LDC endpoints on guest domains. Using this firmware might result in live migration problems, see [“Live Migration Might Fail With Unable to restore ldc resource state on target Domain Migration of LDom failed”](#) on page 50.

**TABLE 1–5** Oracle VM Server for SPARC 3.1.1.1, Oracle VM Server for SPARC 3.1.1, and Oracle VM Server for SPARC 3.1: Required System Firmware Version Patches

Platform Type	3.1.1.1 System Firmware Patch	3.1.1 System Firmware Patch	3.1 System Firmware Patch
SPARC M6-32 server	19525299	18314603	17264114
SPARC M5-32 server	19525299	18314603	17264114
SPARC T5-1B server	19264419	18314586	17264114
SPARC T5-2 server	19264421	18314609	17264122
SPARC T5-4 server	19264423	18314602	17264131
SPARC T5-8 server	19264423	18314602	17264131
Netra SPARC T5-1B server	19264425	18314600	17264110
SPARC T4-1 server	151295-02	150676-05	150676-01
SPARC T4-2 server	151296-02	150677-05	150677-01
SPARC T4-4 server	151297-02	150678-04	150678-01
SPARC T4-1B server	151298-02	150679-04	150679-01
Netra SPARC T4-1 server	151299-02	150680-05	150680-01
Netra SPARC T4-2 server	151300-02	150681-05	150681-01
Netra SPARC T4-1B server	151301-02	150682-04	150682-01

## Minimum Version of Software Required

You can run the Oracle VM Server for SPARC 3.1.1.1, Oracle VM Server for SPARC 3.1.1, or Oracle VM Server for SPARC 3.1 software with at least the minimum versions of the other software components. In such a configuration, you might not have access to all the features of the latest Oracle VM Server for SPARC software. For the best results in production environments, run with the recommended system firmware version described in [“Required](#)

[Software to Enable the Latest Oracle VM Server for SPARC Features](#)” on page 16, and with the Oracle Solaris OS version described in [“Required Oracle Solaris OS Versions”](#) on page 13.

You can apply the latest Oracle VM Server for SPARC package to a system running at least the following versions of software. The minimum software versions are platform specific and depend on the requirements of the CPU in the machine. The minimum Oracle Solaris OS version for a given CPU type applies to all domain types (control, service, I/O, and guest). For information about the minimum versions of the Oracle Solaris OS for supported server platforms, see the data sheet for your server platform at <http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html>.

- Fujitsu M10 systems: XCP2012
- SPARC M6 servers: 9.1.0.g
- SPARC M5 servers: 9.0.1.x
- SPARC T5 servers: 9.0.0.x
- SPARC T4 servers: 8.2.1.b
- SPARC T3 servers: 8.2.1.b
- UltraSPARC T2 Plus servers: 7.4.4.f
- UltraSPARC T2 servers: 7.4.4.f

---

**Note** – On a guest domain, you can run any OS version that is supported by the platform.

---

## Direct I/O Hardware and Software Requirements

To successfully use the direct I/O (DIO) features to assign direct I/O devices to domains, you must run the appropriate software, firmware, or both and use supported PCIe cards.

The dynamic direct I/O feature is supported on Fujitsu M10 systems. The static direct I/O feature is supported on the SPARC T3, SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

- **Hardware Requirements.** Only certain PCIe cards can be used as a direct I/O endpoint device on an I/O domain. You can still use other cards in your Oracle VM Server for SPARC environment but they cannot be used with the DIO feature. Instead, they can be used for service domains and for I/O domains that have entire root complexes assigned to them.

Refer to your platform's hardware documentation to verify which cards can be used on your platform. For an up-to-date list of supported PCIe cards, see <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=REFERENCE&id=1325454.1>.

For an up-to-date list of supported devices on Fujitsu M10 systems, see “Fujitsu M10 systems PCI Card Installation Guide” in the product notes for your model at <http://www.fujitsu.com/global/services/computing/server/sparc/downloads/manual/>

- **Firmware Requirements.** To use the dynamic direct I/O feature on Fujitsu M10 systems, implement the firmware requirements in *Fujitsu M10 Systems Product Notes*.

- **Software Requirements.** To use the DIO feature, the following domains must run the supported OS:
  - **primary domain or a non-primary root domain.** At least the Oracle Solaris 10 9/10 OS plus patch ID 145868-01 or the Oracle Solaris 11 OS.  
The recommended practice is for all domains to run at least the Oracle Solaris 10 1/13 OS plus the required patches in [Table 1–3](#) or the Oracle Solaris 11.1.10.5.0 OS.  
To use the dynamic direct I/O feature on Fujitsu M10 systems, implement the software requirements in *Fujitsu M10 Systems Product Notes*.
  - **I/O domain.** Any Oracle Solaris OS that is supported by the platform.  
To use the dynamic direct I/O feature on Fujitsu M10 systems, implement the software requirements in *Fujitsu M10 Systems Product Notes*.

## PCIe SR-IOV Hardware and Software Requirements

The dynamic PCIe SR-IOV feature is supported on the SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems for Ethernet devices. The static PCIe SR-IOV feature is supported on the SPARC T3, SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

---

**Note** – Before you deploy InfiniBand SR-IOV to your Oracle VM Server for SPARC 3.1 environment, review the information in “[InfiniBand SR-IOV Issues](#)” on [page 48](#).

---

- **Hardware Requirements.**  
Refer to your platform's hardware documentation to verify which cards can be used on your platform. For an up-to-date list of supported PCIe cards, see <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=REFERENCE&id=1325454.1>.
  - **Ethernet SR-IOV.** To use the SR-IOV feature, you can use on-board PCIe SR-IOV devices as well as PCIe SR-IOV plug-in cards. All on-board SR-IOV devices in a given platform are supported unless otherwise explicitly stated in the platform documentation.
  - **InfiniBand SR-IOV.** InfiniBand devices are supported on the SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.
  - **Fibre Channel SR-IOV.** Fibre Channel devices are supported on the SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.  
For an up-to-date list of supported devices on Fujitsu M10 systems, see “Fujitsu M10 systems PCI Card Installation Guide” in the product notes for your model at <http://www.fujitsu.com/global/services/computing/server/sparc/downloads/manual/>
- **Firmware Requirements.**

- **Ethernet SR-IOV.** To use the dynamic SR-IOV feature, SPARC T4 systems must run at least version 8.4.0.a of the system firmware. SPARC T5, SPARC M5, and SPARC M6 platforms must run at least version 9.1.0.a of the system firmware. Fujitsu M10 systems must run at least version XCP2210 of the system firmware. The SPARC T3 platform supports only the static SR-IOV feature. The SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms support only static SR-IOV if not running the most up-to-date system firmware.

To use the SR-IOV feature, PCIe SR-IOV devices must run at least device firmware version 3.01. Perform the following steps to update the firmware for the Sun Dual 10-Gigabit Ethernet SFP+ PCIe 2.0 network adapters:

1. Determine whether you need to upgrade the FCode version on the device.

Perform these commands from the ok prompt:

```
{0} ok cd path-to-device
{0} ok .properties
```

The version value in the output must be one of the following:

```
LP      Sun Dual 10GbE SFP+ PCIe 2.0 LP FCode 3.01 4/2/2012
PEM     Sun Dual 10GbE SFP+ PCIe 2.0 EM FCode 3.01 4/2/2012
FEM     Sun Dual 10GbE SFP+ PCIe 2.0 FEM FCode 3.01 4/2/2012
```

2. Download patch ID 13932765 from [My Oracle Support \(https://support.oracle.com/CSP/ui/flash.html#tab=PatchHomePage\(page=PatchHomePage&id=h0wvdx6\(\)\)\)](https://support.oracle.com/CSP/ui/flash.html#tab=PatchHomePage(page=PatchHomePage&id=h0wvdx6())).
3. Install the patch.

The patch package includes a document that describes how to use the tool to perform the upgrade.

- **InfiniBand SR-IOV.** To use this feature, your system must run at least the following version of the system firmware:
  - **SPARC T4** – 8.4
  - **SPARC T5, SPARC M5, and SPARC M6** – 9.1.0.x
  - **Fujitsu M10 Systems** – XCP2210

To support the Dual 40-Gigabit (4x) InfiniBand Host Channel Adapter M2 as an InfiniBand SR-IOV device, the card or express module must run at least version 2.11.2010 of the firmware. You can obtain this version of the firmware by installing the following patches:

- **Low Profile (X4242A)** – Patch ID 16340059
- **Express Module (X4243A)** – Patch ID 16340042

Use the Oracle Solaris 11.1 `fwflash` command to list and update the firmware in the primary domain. To list the current firmware version, use the `fwflash -lc IB` command. To update the firmware, use the `fwflash -f firmware-file -d device` command. See the [fwflash\(1M\)](#) man page.

To use InfiniBand SR-IOV, ensure that InfiniBand switches have at least firmware version 2.1.2. You can obtain this version of the firmware by installing the following patches:

- **Sun Datacenter InfiniBand Switch 36 (X2821A-Z)** – Patch ID 16221424
- **Sun Network QDR InfiniBand GatewaySwitch (X2826A-Z)** – Patch ID 16221538

For information about how to update the firmware, see your InfiniBand switch documentation.

- **Fibre Channel SR-IOV.** To use this feature, your system must run at least the following version of the system firmware:
  - **SPARC T4** – 8.4.2.c
  - **SPARC T5, SPARC M5, and SPARC M6** – 9.1.2.d
  - **Fujitsu M10 Systems** – XCP2210

The firmware on the Sun Storage 16 Gb Fibre Channel Universal HBA, Emulex must be at least revision 1.1.60.1 to enable the Fibre Channel SR-IOV feature. The installation instructions are provided with the firmware.




---

**Caution** – Only perform the firmware update to the Fibre Channel card if you plan to use the Fibre Channel SR-IOV feature.

---

- **Software Requirements.**
  - **Ethernet SR-IOV.** To use the SR-IOV feature, all domains must be running at least the Oracle Solaris 11.1.10.5.0 OS or the Oracle Solaris 10 1/13 OS plus the required patches in [Table 1–3](#).
  - **InfiniBand SR-IOV.** The following domains must run the supported Oracle Solaris OS:
    - The primary domain or a non-primary root domain must be running at least the Oracle Solaris 11.1.10.6.0 OS.
    - The I/O domains can run at least the Oracle Solaris 11.1.10.6.0 OS or the Oracle Solaris 10 1/13 OS plus the required patches.
    - Update the `/etc/system` file on any root domain that has an InfiniBand SR-IOV physical function from which you plan to configure virtual functions.

```
set ldc:ldc_mactable_entries = 0x20000
```

Update the `/etc/system` file on the I/O domain to which you add a virtual function.

```
set rds3:rds3_fmr_pool_size = 16384
```

- **Fibre Channel SR-IOV.** To use the SR-IOV feature, all domains must be running at least the Oracle Solaris 11.1.17.4.0 OS or the Oracle Solaris 10 1/13 OS plus Oracle VM Server for SPARC 3.1.1 patch ID 150817-02 and the required patches in [Table 1-2](#).

## Non-primary Root Domain Hardware and Software Requirements

Non-primary root domains can be used by the direct I/O and SR-IOV features. This feature is supported on the SPARC T4, SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

- **Hardware Requirements.**

In addition to the PCIe cards for the direct I/O and SR-IOV features in <https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=REFERENCE&id=1325454.1>, other PCIe

cards can be used only in non-primary root domains. To determine which cards you can use on your platform, see your platform's hardware documentation.

- **Firmware Requirements.**

SPARC T4 platforms must run at least version 8.4.0.a of the system firmware.

SPARC T5, SPARC M5, and SPARC M6 platforms must run at least version 9.1.0.x of the system firmware.

Fujitsu M10 systems must run at least version XCP2210 of the system firmware.

- **Software Requirements.**

All domains must be running at least the Oracle Solaris 11.1.10.5.0 OS or the Oracle Solaris 10 1/13 OS plus the required patches in [Table 1-3](#).

## Recovery Mode Hardware and Software Requirements

The recovery mode feature automatically recovers domain configurations that cannot be booted because of faulty or missing resources.

- **Hardware Requirements.**

The recovery mode feature is supported on the SPARC T5, SPARC M5, and SPARC M6 platforms and Fujitsu M10 systems.

- **Firmware Requirements.**

SPARC T5, SPARC M5, and SPARC M6 platforms must run at least version 9.1.0.x of the system firmware.

To use the recovery mode feature on Fujitsu M10 systems, implement the firmware requirements in *Fujitsu M10 Systems Product Notes*.

- **Software Requirements.**

All domains must be running at least the Oracle Solaris 11.1.10.5.0 OS or the Oracle Solaris 10 1/13 OS plus the required patches in [Table 1–3](#).

## Location of the Oracle VM Server for SPARC Software

You can obtain the latest packages for both the Oracle Solaris 10 OS and Oracle Solaris 11 OS for the Oracle VM Server for SPARC 3.1 release. Note that the Oracle VM Server for SPARC software is included by default with the Oracle Solaris 11 OS.

- **Oracle Solaris 10 OS.** Download the `OVM_Server_SPARC-3_1.zip` Oracle VM Server for SPARC package from My Oracle Support. See “[How to Download the Logical Domains Manager Software \(Oracle Solaris 10\)](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.
- **Oracle Solaris 11 OS.** Obtain the `ldomsmanager` package from the Oracle Solaris 11 Support Repository or from My Oracle Support. See “[How to Upgrade to the Oracle VM Server for SPARC 3.1 Software \(Oracle Solaris 11\)](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

Also see the article “[How to Update Oracle Solaris 11 Systems Using Support Repository Updates](http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-018-howto-update-s11-1572261.html)” (<http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-018-howto-update-s11-1572261.html>) and “[How to Update to Oracle Solaris 11.1 Using the Image Packaging System](http://www.oracle.com/technetwork/articles/servers-storage-admin/howto-update-11dot1-ips-1866781.html)” (<http://www.oracle.com/technetwork/articles/servers-storage-admin/howto-update-11dot1-ips-1866781.html>).

For Oracle Solaris 10, the `OVM_Server_SPARC-3_1.zip` file that you download contains the following:

- Oracle VM Server for SPARC software (`SUNWldm.v`)
- The `ldm(1M)`, `ldmconfig(1M)`, and `ldmd(1M)` man pages in the `SUNWldm.v` package, which are installed when the package is installed
- Installation script for Oracle VM Server for SPARC software (`install-ldm`)
- Oracle VM Server for SPARC Management Information Base (`SUNWldmib`)
- Physical-to-Virtual Conversion Tool (`SUNWldmp2v`)

The directory structure of the zip file is similar to the following:

```
Install/
  install-ldm
Product/
  Japanese/
    README.txt
    SUNWjldm.v
    SUNWjldmp2v
  SUNWldm.v
```

SUNWldmib  
 SUNWldmp2v  
 README.txt

## Location of Patches

You can find system firmware for your platform at <http://www.oracle.com/technetwork/systems/patches/firmware/index.html>.

You can find the required the Logical Domains Manager and Oracle Solaris OS patches at <http://support.oracle.com>.

## Location of Documentation

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html>.

For the Oracle Solaris 10 OS, the Oracle VM Server for SPARC man pages are installed on your system as part of the SUNWldm.v and SUNWldmp2v packages. You can install the Japanese translations of the man pages on your Oracle Solaris 10 system with the SUNWjldm.v and SUNWjldmp2v packages.

For the Oracle Solaris 11 OS, the Oracle VM Server for SPARC man pages and the Japanese translations are installed on your system as part of the ldomsmanager package.

The following table shows the documentation that is available for this Oracle VM Server for SPARC release. These documents are available in both HTML and PDF formats unless otherwise indicated.

TABLE 1-6 Related Documentation

Application	Title
Oracle VM Server for SPARC Software	<p><i>Oracle VM Server for SPARC 3.1 Administration Guide</i></p> <p><i>Oracle VM Server for SPARC 3.1 Security Guide</i></p> <p><i>Oracle VM Server for SPARC 3.1 Reference Manual</i></p> <p><i>Oracle VM Server for SPARC 3.1.1 and 3.1 Release Notes</i></p>
Oracle VM Server for SPARC drd(1M) and vntsd(1M) man pages	<p>Oracle Solaris OS Reference Manuals:</p> <ul style="list-style-type: none"> <li>■ Oracle Solaris 10 Documentation (<a href="http://www.oracle.com/technetwork/documentation/solaris-10-192992.html">http://www.oracle.com/technetwork/documentation/solaris-10-192992.html</a>)</li> <li>■ Oracle Solaris 11.1 Documentation (<a href="http://docs.oracle.com/cd/E26502_01">http://docs.oracle.com/cd/E26502_01</a>)</li> </ul>

TABLE 1-6 Related Documentation (Continued)

Application	Title
Oracle Solaris OS: Installation and Configuration	Oracle Solaris OS Installation and Configuration Guides: <ul style="list-style-type: none"> <li>▪ <a href="http://www.oracle.com/technetwork/documentation/solaris-10-192992.html">Oracle Solaris 10 Documentation (http://www.oracle.com/technetwork/documentation/solaris-10-192992.html)</a></li> <li>▪ <a href="http://docs.oracle.com/cd/E26502_01">Oracle Solaris 11.1 Documentation (http://docs.oracle.com/cd/E26502_01)</a></li> </ul>
Oracle VM Server for SPARC and Oracle Solaris OS Security	Oracle VM Server for SPARC White Paper and Oracle Solaris OS Security Guides: <ul style="list-style-type: none"> <li>▪ <a href="http://www.oracle.com/technetwork/articles/systems-hardware-architecture/secure-ovm-sparc-deployment-294062.pdf">Secure Deployment of Oracle VM Server for SPARC (http://www.oracle.com/technetwork/articles/systems-hardware-architecture/secure-ovm-sparc-deployment-294062.pdf)</a></li> <li>▪ <a href="#">Oracle Solaris 10 Security Guidelines</a></li> <li>▪ <a href="#">Oracle Solaris 11 Security Guidelines</a></li> </ul>

You can find documentation that relates to your server, software, or the Oracle Solaris OS at <http://www.oracle.com/technetwork/indexes/documentation/index.html>. Use the Search box to find the documents and the information that you need.

## Related Software

### Software That Can Be Used With the Oracle VM Server for SPARC Software

This section describes the software that is compatible with and can be used with the Oracle VM Server for SPARC software. Be sure to check the software documentation or your platform documentation to find the version number of the software that is available for your version of the Oracle VM Server for SPARC software and your platform.

- **Oracle VM Manager** is a web-based user interface that you can use to manage the Oracle VM environment. For more information about Oracle VM Manager, see the [Oracle VM Documentation \(http://www.oracle.com/technetwork/documentation/vm-096300.html\)](http://www.oracle.com/technetwork/documentation/vm-096300.html).
- **SunVTS** functionality is available in the control domain and guest domains on certain Oracle VM Server for SPARC releases and certain platforms. SunVTS is a validation test suite that provides a comprehensive diagnostic tool that tests and validates Oracle's Sun hardware by verifying the connectivity and proper functioning of most hardware controllers and devices on Oracle's Sun servers. For more information about SunVTS, refer to *SunVTS 7.0 Software*.

- **Explorer Data Collector** can be used with the Oracle VM Server for SPARC software enabled on the control domain. Explorer is a diagnostic data collection tool. The tool comprises shell scripts and a few binary executables. For more information, see *Oracle Explorer User's Guide* (<http://docs.oracle.com/cd/E19957-01/819-6613/819-6613.pdf>).
- **Oracle Solaris Cluster** software can be used in a guest domain with some restrictions. See the Oracle Solaris Cluster documentation for more information about any restrictions and about the Oracle Solaris Cluster software in general.
- **Oracle Enterprise Manager Ops Center** enables you to manage physical and virtual system resources. This solution simplifies resource discovery and monitoring, provides operating system and firmware provisioning, performs comprehensive update and patch management, manages virtual environments such as Oracle Solaris Zones and Oracle VM Server for SPARC, and supports hardware management from power up to production. For more information, see <http://www.oracle.com/us/products/enterprise-manager/index.html>.

## System Controller Software That Interacts With Oracle VM Server for SPARC

The following system controller (SC) software interacts with the Oracle VM Server for SPARC 3.1 software:

- **Sun Integrated Lights Out Manager (ILOM) 3.0** is the system management firmware that you can use to monitor, manage, and configure SPARC T-Series and M-Series systems. ILOM is preinstalled on these platforms and can be used on supported servers with the Oracle VM Server for SPARC 3.1 software enabled. Refer to *Sun Integrated Lights Out Manager 3.0 User's Guide* for features and tasks that are common to Oracle's Sun rack-mounted servers or blade servers that support ILOM. Other user documents present ILOM features and tasks that are specific to the server platform that you are using. You can find the ILOM platform-specific information within the documentation set that accompanies your system.
- **Netra Data Plane Software Suite** is a complete board software package solution. The software provides an optimized rapid development and runtime environment on top of multithread partitioning firmware for Sun CMT platforms. The Logical Domains Manager contains some `ldm` subcommands (`add-vdpcs`, `rm-vdpcs`, `add-vdpcc`, and `rm-vdpcc`) for use with this software. For more information about this software, see the documentation at <http://docs.oracle.com/cd/E19282-01/>.
- **Fujitsu M10 Systems eXtended System Control Facility (XSCF)** is the system management firmware that you can use to monitor, manage, and configure Fujitsu M10 systems. XSCF is preinstalled on these systems and can be used with the Oracle VM Server for SPARC 3.1 software enabled. For more information about this software, see *Fujitsu M10 Systems System Operation and Administration Guide* and *Fujitsu M10 Systems XSCF*

Reference Manual in the Product Notes for your model at <http://www.fujitsu.com/global/services/computing/server/sparc/downloads/manual/>.

## Optional Software

Oracle VM Server for SPARC Management Information Base (MIB) software can help you enable third-party applications to perform remote monitoring and a few control operations. For more information, see [Chapter 17, “Using the Oracle VM Server for SPARC Management Information Base Software,”](#) in *Oracle VM Server for SPARC 3.1 Administration Guide*.

## Upgrading to the Current Oracle VM Server for SPARC Software



**Caution** – Do not downgrade to older versions of individual software and firmware components. Such downgrades are not recommended and might lead to unexpected behavior and failures.

### Upgrading to the Oracle VM Server for SPARC 3.1.1.1 Software

On a control domain that runs the Oracle Solaris 10 OS, you can upgrade the Oracle VM Server for SPARC software to the Oracle VM Server for SPARC 3.1.1.1 software. For example, you can upgrade directly to Oracle VM Server for SPARC 3.1.1.1 from the 3.1 or 3.1.1 versions by applying the Oracle VM Server for SPARC 3.1.1.1 patch.

- **Oracle Solaris 11 OS.** Install the Oracle VM Server for SPARC 3.1.1.1 release on your control domain. See [“How to Upgrade to the Oracle VM Server for SPARC 3.1 Software \(Oracle Solaris 11\)”](#) in *Oracle VM Server for SPARC 3.1 Administration Guide*.
- **Oracle Solaris 10 OS.** Install the Oracle VM Server for SPARC 3.1 release on your control domain, if necessary, and apply the Oracle VM Server for SPARC 3.1.1.1 patch.
  - **Oracle VM Server for SPARC 3.1.1.** Disable the `ldmd` daemon, apply Oracle VM Server for SPARC 3.1.1.1 patch 150817-03, and re-enable the daemon.

```
primary# svcadm -v disable -st ldmd
primary# patchadd 150817-03
primary# svcadm -v enable ldmd
```

- **Oracle VM Server for SPARC 3.1.** Disable the `ldmd` daemon, apply Oracle VM Server for SPARC 3.1.1.1 patch 150817-03, and re-enable the daemon.

```
primary# svcadm -v disable -st ldmd
primary# patchadd 150817-03
primary# svcadm -v enable ldmd
```

- **Older Oracle VM Server for SPARC releases.** Install the Oracle VM Server for SPARC 3.1 software. See “[How to Upgrade to the Oracle VM Server for SPARC 3.1 Software \(Oracle Solaris 10\)](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

Then, disable the `ldmd` daemon, apply Oracle VM Server for SPARC 3.1.1.1 patch 150817-03, and re-enable the daemon.

```
primary# svcadm -v disable -st ldmd
primary# patchadd 150817-03
primary# svcadm -v enable ldmd
```

---

**Note** – Logical Domains Manager 3.1.1.1 is not supported on UltraSPARC T1 systems.

---

## Upgrading to the Oracle VM Server for SPARC 3.1.1 Software

You can upgrade a control domain running an older version of the Oracle VM Server for SPARC software to the Oracle VM Server for SPARC 3.1.1 software.

- **Oracle Solaris 11 OS.** Install the Oracle VM Server for SPARC 3.1.1 release on your control domain. See “[How to Upgrade to the Oracle VM Server for SPARC 3.1 Software \(Oracle Solaris 11\)](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.
- **Oracle Solaris 10 OS.** Install the Oracle VM Server for SPARC 3.1 release on your control domain, if necessary, and apply the Oracle VM Server for SPARC 3.1.1 patch.
  - **Oracle VM Server for SPARC 3.1.** Disable the `ldmd` daemon, apply Oracle VM Server for SPARC 3.1.1 patch 150817-02, and re-enable the daemon.

```
primary# svcadm -v disable -st ldmd
primary# patchadd 150817-02
primary# svcadm -v enable ldmd
```

- **Older Oracle VM Server for SPARC releases.** Install the Oracle VM Server for SPARC 3.1 software. See “[How to Upgrade to the Oracle VM Server for SPARC 3.1 Software \(Oracle Solaris 10\)](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

Then, disable the `ldmd` daemon, apply Oracle VM Server for SPARC 3.1.1 patch 150817-02, and re-enable the daemon.

```
primary# svcadm -v disable -st ldmd
primary# patchadd 150817-02
primary# svcadm -v enable ldmd
```

---

**Note** – Logical Domains Manager 3.1.1 is not supported on UltraSPARC T1 systems.

---

## Upgrading to the Oracle VM Server for SPARC 3.1 Software

You can upgrade to the Oracle VM Server for SPARC 3.1 software for either Oracle Solaris 10 or Oracle Solaris 11:

- **Oracle Solaris 10:** See “[How to Upgrade to the Oracle VM Server for SPARC 3.1 Software \(Oracle Solaris 10\)](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.
- **Oracle Solaris 11:** Use the `pkg update` command to specify one of the Oracle Solaris 11 releases that includes the Oracle VM Server for SPARC 3.1 software (Oracle Solaris 11.1.10 through Oracle Solaris 11.1.16). When no release is specified, the `pkg update` command installs the latest SRU, which includes the Oracle VM Server for SPARC 3.1.1 software. See “[How to Upgrade to the Oracle VM Server for SPARC 3.1 Software \(Oracle Solaris 11\)](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

---

**Note** – Logical Domains Manager 3.1 is not supported on UltraSPARC T1 systems.

---

## Deprecated Oracle VM Server for SPARC Features

The following Oracle VM Server for SPARC features are deprecated in this release of the software and might be removed from the product at any time:

- The network interface functionality of an Oracle VM Server for SPARC vsw driver was deprecated in Oracle Solaris 11.1. The Oracle VM Server for SPARC vsw driver continues to provide virtual network switching functionality for guest domains. See “[Oracle Solaris 11 Networking Overview](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

---

**Note** – The network interface functionality continues to be supported for Oracle Solaris 10 service domains.

---

- Setting the `threading` property to `max-ipc` to manage single-thread workloads on SPARC T4 platforms is deprecated. By default, domains are created for maximum throughput, and the Oracle Solaris OS automatically uses the critical thread API to optimize for single-thread workloads.
- Using the `-c` option of the `add-vcpu`, `set-vcpu`, and `rm-vcpu` subcommands to manage hard partitioning is deprecated. Instead, use the `add-core`, `set-core`, or `rm-core` subcommand to assign whole cores. Also, use the `add-domain` or `set-domain` subcommand to specify the CPU cap (`max-cores`).
- Using the Hybrid I/O feature is deprecated in favor of using the single-root I/O virtualization (SR-IOV) feature.

- Migrations of domains that run an Oracle Solaris OS older than Oracle Solaris 10 9/10 are known as “uncooperative” migrations. The live migration feature has deprecated the use of the uncooperative migration feature. In the future, these types of migrations might be explicitly rejected. Therefore, ensure that the guest domain that you want to migrate runs at least the Oracle Solaris 10 9/10 OS or Oracle Solaris 11 OS.

## Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 3.1.1.1, Oracle VM Server for SPARC 3.1.1, and Oracle VM Server for SPARC 3.1 software.

### General Issues

This section describes general known issues about this release of the Oracle VM Server for SPARC software that are broader than a specific bug number. Workarounds are provided where available.

#### **Cannot Unbind Domains When They Provide Services to Each Other**

Do not create a circular dependency between two domains in which each domain provides services to the other. Such a configuration creates a single point of failure condition where an outage in one domain causes the other domain to become unavailable. Circular dependency configurations also prevent you from unbinding the domains after they have been bound initially.

The Logical Domains Manager does not prevent the creation of circular domain dependencies.

If the domains cannot be unbound due to a circular dependency, remove the devices that cause the dependency and then attempt to unbind the domains.

#### **Guest Domain Cannot Run the Oracle Solaris 10 OS When More Than 1024 CPUs Are Assigned**

A guest domain that has been assigned more than 1024 CPUs cannot run the Oracle Solaris 10 OS. In addition, you cannot use CPU DR to shrink the number of CPUs below 1024 to run the Oracle Solaris 10 OS.

To work around this problem, unbind the guest domain, remove CPUs until you have no more than 1024 CPUs, and then rebind the guest domain. You can then run the Oracle Solaris 10 OS on this guest domain.

## Avoid Creating a Configuration Where Two Domains Provide Services to Each Other

Avoid creating a configuration where two domains provide services to each other. In such a case, an outage in one domain will take down the other domain. In addition, such domains cannot be unbound if they are bound with such a configuration. The Logical Domains Manager currently does not block such circular dependencies.

If you cannot unbind a domain because of this sort of dependency, remove the devices that cause the circular dependency and then attempt the unbind again.

## Upgrading From Oracle Solaris 10 OS Older Than Oracle Solaris 10 5/08 OS

If the control domain is upgraded from an Oracle Solaris 10 OS version older than Oracle Solaris 10 5/08 OS (or without patch 127127-11), and if volume manager volumes were exported as virtual disks, the virtual disk back ends must be re-exported with `options=slice` after the Logical Domains Manager has been upgraded. See “[Exporting Volumes and Backward Compatibility](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

## Service Processor and System Controller Are Interchangeable Terms

For discussions in Oracle VM Server for SPARC documentation, the terms service processor (SP) and system controller (SC) are interchangeable.

## In Certain Conditions, a Guest Domain's Solaris Volume Manager Configuration or Metadevices Can Be Lost

If a service domain is running a version of Oracle Solaris 10 OS prior to Oracle Solaris 10 1/13 OS and is exporting a physical disk slice as a virtual disk to a guest domain, then this virtual disk will appear in the guest domain with an inappropriate device ID. If that service domain is then upgraded to Oracle Solaris 10 1/13 OS, the physical disk slice exported as a virtual disk will appear in the guest domain with no device ID.

This removal of the device ID of the virtual disk can cause problems to applications attempting to reference the device ID of virtual disks. In particular, Solaris Volume Manager might be unable to find its configuration or to access its metadevices.

**Workaround:** After upgrading a service domain to Oracle Solaris 10 1/13 OS, if a guest domain is unable to find its Solaris Volume Manager configuration or its metadevices, perform the following procedure.

### ▼ How to Find a Guest Domain's Solaris Volume Manager Configuration or Metadevices

- 1 Boot the guest domain.

- 2 Disable the `devid` feature of Solaris Volume Manager by adding the following lines to the `/kernel/drv/md.conf` file:**

```
md_devid_destroy=1;
md_keep_repl_state=1;
```

- 3 Reboot the guest domain.**

After the domain has booted, the Solaris Volume Manager configuration and metadevices should be available.

- 4 Check the Solaris Volume Manager configuration and ensure that it is correct.**

- 5 Re-enable the Solaris Volume Manager `devid` feature by removing from the `/kernel/drv/md.conf` file the two lines that you added in Step 2.**

- 6 Reboot the guest domain.**

During the reboot, you will see messages similar to this:

```
NOTICE: mddb: unable to get devid for 'vdc', 0x10
```

These messages are normal and do not report any problems.

## Memory Size Requirements

The Oracle VM Server for SPARC software does not impose a memory size limitation when you create a domain. The memory size requirement is a characteristic of the guest operating system. Some Oracle VM Server for SPARC functionality might not work if the amount of memory present is smaller than the recommended size. For recommended and minimum memory requirements for the Oracle Solaris 10 OS, see [“System Requirements and Recommendations” in Oracle Solaris 10 8/11 Installation Guide: Planning for Installation and Upgrade](#). For recommended and minimum memory requirements for the Oracle Solaris 11 OS, see [Oracle Solaris 11 Release Notes](#) and [Oracle Solaris 11.1 Release Notes](#).

The OpenBoot PROM has a minimum size restriction for a domain. Currently, that restriction is 12 Mbytes. If you have a domain smaller than that size, the Logical Domains Manager will automatically boost the size of the domain to 12 Mbytes. The minimum size restriction for a Fujitsu M10 system is 256 Mbytes. Refer to the release notes for your system firmware for information about memory size requirements.

The memory dynamic reconfiguration (DR) feature enforces 256-Mbyte alignment on the address and size of the memory involved in a given operation. See [“Memory Alignment” in Oracle VM Server for SPARC 3.1 Administration Guide](#).

## Booting a Large Number of Domains

You can boot the following number of domains depending on your platform:

- Up to 256 on Fujitsu M10 systems per physical partition
- Up to 128 on SPARC M6 systems per physical domain
- Up to 128 on SPARC M5 systems per physical domain
- Up to 128 on SPARC T5 systems
- Up to 128 on SPARC T4 servers
- Up to 128 on SPARC T3 servers
- Up to 128 on UltraSPARC T2 Plus servers
- Up to 64 on UltraSPARC T2 servers

If unallocated virtual CPUs are available, assign them to the service domain to help process the virtual I/O requests. Allocate 4 to 8 virtual CPUs to the service domain when creating more than 32 domains. In cases where maximum domain configurations have only a single CPU in the service domain, do not put unnecessary stress on the single CPU when configuring and using the domain. The virtual switch (vsw) services should be spread across all the network adapters available in the machine. For example, if booting 128 domains on a Sun SPARC Enterprise T5240 server, create 4 vsw services, each serving 32 virtual net (vnet) instances. Assigning more than 32 vnet instances per vsw service could cause hard hangs in the service domain.

To run the maximum configurations, a machine needs an adequate amount of memory to support the guest domains. The amount of memory is dependent on your platform and your OS. See the documentation for your platform, [Oracle Solaris 10 8/11 Installation Guide: Planning for Installation and Upgrade](#), [Installing Oracle Solaris 11 Systems](#), and [Installing Oracle Solaris 11.1 Systems](#).

Memory and swap space usage increases in a guest domain when the vsw services used by the domain provide services to many virtual networks in multiple domains. This increase is due to the peer-to-peer links between all the vnet instances connected to the vsw. The service domain benefits from having extra memory. The recommended minimum is four Gbytes when running more than 64 domains. Start domains in groups of 10 or fewer and wait for them to boot before starting the next batch. The same advice applies to installing operating systems on domains. You can reduce the number of links by disabling inter-vnet links. See “[Inter-Vnet LDC Channels](#)” in [Oracle VM Server for SPARC 3.1 Administration Guide](#).

## Cleanly Shutting Down and Power Cycling an Oracle VM Server for SPARC System

If you have made any configuration changes since last saving a configuration to the SC, before you attempt to power off or power cycle an Oracle VM Server for SPARC system, make sure that you save the latest configuration that you want to keep.

## ▼ How to Power Off a System With Multiple Active Domains

- 1 Shut down, stop, and unbind all the non-I/O domains.
- 2 Shut down, stop, and unbind any active I/O domains.
- 3 Halt the primary domain.

Because no other domains are bound, the firmware automatically powers off the system.

## ▼ How to Power Cycle the System

- 1 Shut down, stop, and unbind all the non-I/O domains.
- 2 Shut down, stop, and unbind any active I/O domains.
- 3 Reboot the primary domain.

Because no other domains are bound, the firmware automatically power cycles the system before rebooting it. When the system restarts, it boots into the domain configuration last saved or explicitly set.

## Memory Size Requested Might Be Different From Memory Allocated

Under certain circumstances, the Logical Domains Manager rounds up the requested memory allocation to either the next largest 8-Kbyte or 4-Mbyte multiple. The following example shows sample output of the `ldm list-domain -l` command, where the constraint value is smaller than the actual allocated size:

```
Memory:
      Constraints: 1965 M
      raddr      paddr5      size
      0x1000000  0x291000000  1968M
```

## Logical Domains Variable Persistence

Variable updates persist across a reboot but not across a power cycle unless the variable updates are either initiated from OpenBoot firmware on the control domain or followed by saving the configuration to the SC.

Note the following conditions:

- When the control domain reboots, if there are no bound guest domains and no delayed reconfiguration in progress, the SC performs a power cycle of the system.
- When the control domain reboots, if guest domains are bound or active (or the control domain is in the middle of a delayed reconfiguration), the SC does not perform a power cycle of the system.

Logical Domains variables for a domain can be specified using any of the following methods:

- At the OpenBoot prompt.
- Using the Oracle Solaris OS `eeeprom(1M)` command.
- Using the Logical Domains Manager CLI (`ldm`).
- In a limited fashion, from the system controller (SC) using the `bootmode` command. This method can be used for only certain variables, and only when in the `factory-default` configuration.

Variable updates that are made by using any of these methods should always persist across reboots of the domain. The variable updates also always apply to any subsequent domain configurations that were saved to the SC.

In Oracle VM Server for SPARC 3.1 software, variable updates do not persist as expected in a few cases:

- All methods of updating a variable persist across reboots of that domain. However, they do not persist across a power cycle of the system unless a subsequent logical domain configuration is saved to the SC.

However in the control domain, updates made using either OpenBoot firmware commands or the `eeeprom` command *do* persist across a power cycle of the system even without subsequently saving a new logical domain configuration to the SC. The `eeeprom` command supports this behavior on SPARC T5, SPARC M5, and SPARC M6 systems, and on SPARC T3 and SPARC T4 systems that run at least version 8.2.1 of the system firmware.

- In all cases, when reverting to the factory-default configuration from a configuration generated by the Logical Domains Manager, all Logical Domains variables start with their default values.

If you are concerned about Logical Domains variable changes, do one of the following:

- Bring the system to the ok prompt and update the variables.
- Update the variables while the Logical Domains Manager is disabled:

```
# svcadm disable ldmd
update variables
# svcadm enable ldmd
```

- When running Live Upgrade, perform the following steps:

```
# svcadm disable -t ldmd
# luactivate be3
# init 6
```

If you modify the time or date on a logical domain, for example, using the `ntpdate` command, the change persists across reboots of the domain but not across a power cycle of the host. To ensure that time changes persist, save the configuration with the time change to the SP and boot from that configuration.

The following Bug IDs have been filed to resolve these issues: 15375997, 15387338, 15387606, and 15415199.

## Oracle's Sun SNMP Management Agent Does Not Support Multiple Domains

Sun Simple Network Management Protocol (SNMP) Management Agent does not support multiple domains. Only a single global domain is supported.

## Delayed Reconfiguration

When the primary domain is in a delayed reconfiguration state, resources that are managed by Oracle VM Server for SPARC are power-managed *only* after the primary domain reboots. Resources that are managed directly by the OS, such as CPUs that are managed by the Solaris Power Aware Dispatcher, are not affected by this state.

## Cryptographic Units

Discrete cryptographic units are present only on UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems.

Cryptographic unit dynamic reconfiguration (DR) enables you to add and remove cryptographic units from a domain. The Logical Domains Manager automatically detects whether a domain allows cryptographic unit DR, and enables the functionality only for those domains. In addition, CPU DR is no longer disabled in domains that have cryptographic units bound and then are running an appropriate version of the Oracle Solaris OS.

## ldmp2v convert Command: VxVM Warning Messages During Boot

Running Veritas Volume Manager (VxVM) 5.x on the Oracle Solaris 10 OS is the only supported (tested) version for the Oracle VM Server for SPARC P2V tool. Older versions of VxVM, such as 3.x and 4.x running on the Solaris 8 and Solaris 9 operating systems, might also work. In those cases, the first boot after running the `ldmp2v convert` command might show warning messages from the VxVM drivers. You can ignore these messages. You can remove the old `VRTS*` packages after the guest domain has booted.

```
Boot device: disk0:a File and args:
SunOS Release 5.10 Version Generic_139555-08 64-bit
Copyright 1983-2009 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Hostname: normal
Configuring devices.
/kernel/drv/sparcv9/vxdmp: undefined symbol 'romp'
WARNING: mod_load: cannot load module 'vxdmp'
WARNING: vxdmp: unable to resolve dependency, module 'misc/ted' not found
/kernel/drv/sparcv9/vxdmp: undefined symbol 'romp'
WARNING: mod_load: cannot load module 'vxdmp'
WARNING: vxdmp: unable to resolve dependency, module 'misc/ted' not found
```

```

/kernel/drv/sparcv9/vxio: undefined symbol 'romp'
WARNING: mod_load: cannot load module 'vxio'
WARNING: vxio: unable to resolve dependency, module 'drv/vxdmp' not found
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
NOTICE: VxVM not started

```

## Oracle Hard Partitioning Requirements for Software Licenses

For information about Oracle's hard partitioning requirements for software licenses, see [Partitioning: Server/Hardware Partitioning \(http://www.oracle.com/us/corporate/pricing/partitioning-070609.pdf\)](http://www.oracle.com/us/corporate/pricing/partitioning-070609.pdf).

## Upgrade Option Not Presented When Using `ldmp2v prepare -R`

The Oracle Solaris Installer does not present the Upgrade option when the partition tag of the slice that holds the root (`/`) file system is not set to `root`. This situation occurs if the tag is not explicitly set when labeling the guest's boot disk. You can use the `format` command to set the partition tag as follows:

```

AVAILABLE DISK SELECTIONS:
0. c0d0 <SUN-DiskImage-10GB cyl 282 alt 2 hd 96 sec 768>
   /virtual-devices@100/channel-devices@200/disk@0
1. c4t2d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
   /pci@400/pci@0/pci@1/scsi@0/sd@2,0
2. c4t3d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
   /pci@400/pci@0/pci@1/scsi@0/sd@3,0
Specify disk (enter its number)[0]: 0
selecting c0d0
[disk formatted, no defect list found]
format> p

```

```

PARTITION MENU:
0      - change '0' partition
1      - change '1' partition
2      - change '2' partition
3      - change '3' partition
4      - change '4' partition
5      - change '5' partition
6      - change '6' partition
7      - change '7' partition
select - select a predefined table
modify - modify a predefined partition table

```

```

name - name the current table
print - display the current table
label - write partition map and label to the disk
!<cmd> - execute <cmd>, then return
quit

partition> 0
Part      Tag      Flag      Cylinders      Size      Blocks
0 unassigned  wm        0              0          (0/0/0)      0

Enter partition id tag[unassigned]: root
Enter partition permission flags[wm]:
Enter new starting cyl[0]: 0
Enter partition size[0b, 0c, 0e, 0.00mb, 0.00gb]: 8g
partition> label
Ready to label disk, continue? y

partition>

```

## Sometimes a Block of Dynamically Added Memory Can be Dynamically Removed Only as a Whole

Due to the way in which the Oracle Solaris OS handles the metadata for managing dynamically added memory, you might later be able to remove only the entire block of memory that was previously dynamically added rather than a proper subset of that memory.

This situation could occur if a domain with a small memory size is dynamically grown to a much larger size, as shown in the following example.

```

primary# ldm list ldom1
NAME STATE FLAGS  CONS VCPU MEMORY UTIL UPTIME
ldom1 active -n-- 5000 2 2G 0.4% 23h

primary# ldm add-mem 16G ldom1

primary# ldm rm-mem 8G ldom1
Memory removal failed because all of the memory is in use.

primary# ldm rm-mem 16G ldom1

primary# ldm list ldom1
NAME STATE FLAGS  CONS VCPU MEMORY UTIL UPTIME
ldom1 active -n-- 5000 2 2G 0.4% 23h

```

**Workaround:** Use the `ldm add-mem` command to sequentially add memory in smaller chunks rather than in chunks larger than you might want to remove in the future.

**Recovery:** Perform one of the following actions:

- Stop the domain, remove the memory, and then restart the domain.
- Reboot the domain, which causes the Oracle Solaris OS to re-allocate its memory management metadata such that the previously added memory can now be removed dynamically in smaller chunks.

## **ldmp2v Command: ufsdump Archiving Method Is No Longer Used**

Restoring `ufsdump` archives on a virtual disk that is backed by a file on a UFS file system might cause the system to hang. In such a case, the `ldmp2v prepare` command will exit. You might encounter this problem when you manually restore `ufsdump` archives in preparation for the `ldmp2v prepare -R /alt root` command when the virtual disk is a file on a UFS file system. For compatibility with previously created `ufsdump` archives, you can still use the `ldmp2v prepare` command to restore `ufsdump` archives on virtual disks that are not backed by a file on a UFS file system. However, the use of `ufsdump` archives is not recommended.

## **Only One CPU Configuration Operation Is Permitted to Be Performed During a Delayed Reconfiguration**

Do not attempt to perform more than one CPU configuration operation on the primary domain while it is in a delayed reconfiguration. If you attempt more CPU configuration requests, they will be rejected.

**Workaround:** Perform one of the following actions:

- Cancel the delayed reconfiguration, start another one, and request the configuration changes that were lost from the previous delayed reconfiguration.
- Reboot the control domain with the incorrect CPU count and then make the allocation corrections after the domain reboots.

## **Oracle VM Server for SPARC 3.1 ldmd Daemon Does Not Start If Multiple Virtual Switches Are Assigned to a Single Network Adapter**

The Oracle VM Server for SPARC 3.0 software inadvertently exposed a capability to assign multiple virtual switches to a single network adapter. This capability is intended only to be used in a specific way by the Oracle VM Manager software.

The Oracle VM Server for SPARC 3.1 software restored the original behavior, which prevents you from assigning multiple virtual switches to a single network adapter. However, if you configured your Oracle VM Server for SPARC 3.0 system to assign multiple virtual switches to a single network adapter, the `ldmd` daemon does not start when you upgrade to Oracle VM Server for SPARC 3.1.

**Workaround:** Perform the following steps:

1. Temporarily re-enable this capability on your Oracle VM Server for SPARC 3.1 system to enable the `ldmd` daemon to start.

```
# svccfg -s ldoms/ldmd setprop ldmd/ovm_manager=true
# svcadm refresh ldmd
# svcadm disable ldmd
# svcadm enable ldmd
```

2. Update your configuration to assign only one virtual switch to a network device.

3. Disable this capability on your Oracle VM Server for SPARC 3.1 system.

```
# svccfg -s ldoms/ldmd setprop ldmd/ovm_manager=false
# svcadm refresh ldmd
# svcadm disable ldmd
# svcadm enable ldmd
```

It is important that you set the `ovm_manager` property to `false` because this property might introduce other side effects in future Oracle VM Server for SPARC releases.

## Oracle Solaris Boot Disk Compatibility

Historically, the Oracle Solaris OS has been installed on a boot disk configured with an SMI VTOC disk label. Starting with the Oracle Solaris 11.1 OS, the OS is installed on a boot disk that is configured with an extensible firmware interface (EFI) GUID partition table (GPT) disk label by default. If the firmware does not support EFI, the disk is configured with an SMI VTOC disk label instead. This situation applies only to SPARC T4 servers that run at least system firmware version 8.4.0, to SPARC T5, SPARC M5, or SPARC M6 servers that run at least system firmware version 9.1.0, and to Fujitsu M10 systems that run at least XCP2230.

The following servers cannot boot from a disk that has an EFI GPT disk label:

- UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 servers no matter which system firmware version is used
- SPARC T4 servers that run system firmware versions prior to 8.4.0
- SPARC T5, SPARC M5, and SPARC M6 servers that run system firmware versions prior to 9.1.0
- Fujitsu M10 systems that run XCP versions prior to 2230

So, an Oracle Solaris 11.1 boot disk that is created on an up-to-date SPARC T4, SPARC T5, SPARC M5, SPARC M6, or Fujitsu M10 system cannot be used on older servers or on servers that run older firmware.

This limitation restrains the ability to use either cold or live migration to move a domain from a recent server to an older server. This limitation also prevents you from using an EFI GPT boot disk image on an older server.

To determine whether an Oracle Solaris 11.1 boot disk is compatible with your server and its firmware, ensure that the Oracle Solaris 11.1 OS is installed on a disk that is configured with an SMI VTOC disk label.

To maintain backward compatibility with systems that run older firmware, use one of the following procedures. Otherwise, the boot disk uses the EFI GPT disk label by default. These procedures show how to ensure that the Oracle Solaris 11.1 OS is installed on a boot disk with an SMI VTOC disk label on a SPARC T4 server with at least system firmware version 8.4.0, on a SPARC T5, SPARC M5, or SPARC M6 server with at least system firmware version 9.1.0, and on a Fujitsu M10 system with at least XCP version 2230.

- **Solution 1:** Remove the `gpt` property so that the firmware does not report that it supports EFI.

1. From the OpenBoot PROM prompt, disable automatic booting and reset the system to be installed.

```
ok setenv auto-boot? false
ok reset-all
```

After the system resets, it returns to the `ok` prompt.

2. Change to the `/packages/disk-label` directory and remove the `gpt` property.

```
ok cd /packages/disk-label
ok "gpt" delete-property
```

3. Begin the Oracle Solaris 11.1 OS installation.

For example, perform a network installation:

```
ok boot net - install
```

- **Solution 2:** Use the `format -e` command to write an SMI VTOC label on the disk to be installed with the Oracle Solaris 11.1 OS.

1. Write an SMI VTOC label on the disk.

For example, select the `label` option and specify the SMI label:

```
# format -e c1d0
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 0
```

2. Configure the disk with a slice 0 and slice 2 that cover the entire disk.

The disk should have no other partitions. For example:

```
format> partition
```

```
partition> print
Current partition table (unnamed):
Total disk cylinders available: 14087 + 2 (reserved cylinders)
```

Part	Tag	Flag	Cylinders	Size	Blocks
0	root	wm	0 - 14086	136.71GB	(14087/0/0) 286698624
1	unassigned	wu	0	0	(0/0/0) 0
2	backup	wu	0 - 14086	136.71GB	(14087/0/0) 286698624
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0

```

        6 unassigned   wm      0          0          (0/0/0)      0
        7 unassigned   wm      0          0          (0/0/0)      0

```

3. Re-write the SMI VTOC disk label.

```

partition> label
[0] SMI Label
[1] EFI Label
Specify Label type[0]: 0
Ready to label disk, continue? y

```

4. Configure your Oracle Solaris Automatic Installer (AI) to install the Oracle Solaris OS on slice 0 of the boot disk.

Change the <disk> excerpt in the AI manifest as follows:

```

<target>
  <disk whole_disk="true">
    <disk_keyword key="boot_disk"/>
    <slice name="0" in_zpool="rpool"/>
  </disk>
  [...]
</target>

```

5. Perform the installation of the Oracle Solaris 11.1 OS.

## Domain Migration Restrictions

The following sections describe restrictions for domain migration. The Logical Domains Manager software and the system firmware versions must be compatible to permit migrations. Also, you must meet certain CPU requirements to ensure a successful domain migration.

Live migration is not qualified and supported on all combinations of the source and target platforms and system firmware versions. For those combinations that cannot perform a live migration, you can perform a cold migration instead.

### Version Restrictions for Migration

This section describes version restrictions for performing live migrations.

- Logical Domains Manager version.** You can perform a live migration in either direction when one system runs the latest version of the Logical Domains Manager and the other system runs at least the immediately preceding version of the Logical Domains Manager. In addition, in the case of the Oracle VM Server for SPARC 3.1.x software, you can perform the live migration of a domain to or from a system that runs version 3.1.x of the Logical Domains Manager to or from a system that runs version 3.0.0.x of the Logical Domains Manager.
- System firmware version.** In general, you can perform a live migration between two systems when both the source and target machines support the appropriate minimum system firmware versions.

The following list shows the platforms that support live migration and the associated minimum system firmware version:

- **UltraSPARC T2 and UltraSPARC T2 Plus platforms** – Version 7.4.5
- **SPARC T3 and SPARC T4 platforms** – Version 8.2.2.c
- **SPARC T5, SPARC M5, and SPARC M6 platforms** – All system firmware versions
- **Fujitsu M10 systems** – All XCP versions

However, some specific platform and firmware combinations do not support live migration. Attempting to perform the live migration of a domain from a system that runs at least system firmware version 8.4 or XCP2210 to a system that runs an older system firmware version fails. The failure occurs because of a hypervisor API mismatch between the newer and older system firmware versions. In this instance, the following message is issued:

```
primary# ldm migrate ldg1 root@target-name
Target Password:
Domain ldg1 is using features of the system firmware that are not supported in
the version of the firmware running on the target machine.
Domain Migration of LDom ldg1 failed
```

Note that you can perform the live migration of a domain from a system that runs system firmware version 8.3 to a system that runs at least system firmware version 8.4 unless the target machine is a SPARC M5-32 system. For more information, see [“Domain Migrations From SPARC T4 Systems That Run System Firmware 8.3 to SPARC T5, SPARC M5, or SPARC M6 Systems Are Erroneously Permitted”](#) on page 62.

System firmware versions 8.4, 9.1, and XCP2230 introduced support for EFI GPT disk labels. By default, virtual disks that are installed when running at least the Oracle Solaris 11.1 OS on those systems have an EFI GPT disk label. You cannot read this disk label on older versions of firmware (such as 9.0.x, 8.3, 7.x, or XCP2221). This situation precludes you from performing a live migration or a cold migration to a system that runs a system firmware version without EFI GPT support. Note that a cold migration also fails in this situation, which is different than the previous limitations.

To determine whether your virtual disk has an EFI GPT disk label, run the `devinfo -i` command on the raw device. The following examples show whether the virtual disk has an SMI VTOC or an EFI GPT disk label:

- **SMI VTOC disk label.** When your virtual disk has an SMI VTOC, you can perform a migration to firmware regardless of whether it supports EFI.

This example indicates that the device has a VTOC label because the `devinfo -i` command reports device-specific information:

```
# devinfo -i /dev/rdisk/c2d0s2
/dev/rdisk/c2d0s2      0      0      73728  512    2
```

- **EFI GPT disk label.** When your virtual disk has an EFI GPT disk label, you can perform a migration only to firmware that has EFI support.

This example indicates that the device has an EFI GPT disk label because the `devinfo -i` command reports an error:

```
# devinfo -i /dev/rdisk/c1d0s0
devinfo: /dev/rdisk/c1d0s0: This operation is not supported on EFI
labeled devices
```

## CPU Restrictions for Migration

If the domain to be migrated is running an Oracle Solaris OS version older than the Oracle Solaris 10 1/13 OS, you might see the following message during the migration:

```
Domain domain-name is not running an operating system that is
compatible with the latest migration functionality.
```

The following CPU requirements and restrictions apply when you run an OS prior to the Oracle Solaris 10 1/13 OS:

- Full cores must be allocated to the migrated domain. If the number of threads in the domain to be migrated is less than a full core, the extra threads are unavailable to any domain until after the migrated domain is rebooted.
- After a migration, CPU dynamic reconfiguration (DR) is disabled for the migrated domain until it has been rebooted. At that time, you can use CPU DR on the migrated domain.
- The target machine must have enough free full cores to provide the number of threads that are required for the migrated domain. After the migration, if a full core is only partially used by the migrated domain, any extra threads are unavailable to any domain until after the migrated domain is rebooted.

These restrictions also apply when you attempt to migrate a domain that is running in OpenBoot or in the kernel debugger. See [“Migrating a Domain From the OpenBoot PROM or a Domain That Is Running in the Kernel Debugger”](#) in *Oracle VM Server for SPARC 3.1 Administration Guide*.

## Version Restrictions for Cross-CPU Migration

You cannot perform live migrations between an UltraSPARC T2, UltraSPARC T2 Plus, or SPARC T3 system and a SPARC T5, SPARC M5, or SPARC M6 system.

You can only perform a live migration between a SPARC T4 system and a SPARC T5, SPARC M5, or SPARC M6 system only if the following requirements are met:

- The SPARC T4 system must run system firmware version 8.4
- The SPARC T5, SPARC M5, or SPARC M6 system must run system firmware version 9.1
- Both the source and target machines must run the Oracle VM Server for SPARC 3.1 software

## Domains That Have Only One Virtual CPU Assigned Might Panic During a Live Migration

**Bug ID 17285751:** Migrating a domain that has only one virtual CPU assigned to it might cause a panic on the guest domain in the function `pg_cmt_cpu_fini()`.

**Workaround:** Assign at least two virtual CPUs to the guest domain before you perform the live migration. For example, use the `ldm add-vcpu number-of-virtual-CPUs domain` command to increase the number of virtual CPUs assigned to the guest domain.

## Oracle VM Server for SPARC MIB Issues

This section summarizes the issues that you might encounter when using Oracle VM Server for SPARC Management Information Base (MIB) software.

### snmpTable Command Does Not Work With the Version 2 or Version 3 Option

**Bug ID 15376861:** You receive empty SNMP tables if you query the Oracle VM Server for SPARC MIB 2.1 software using the `snmpTable` command with the `-v2c` or `-v3` option. The `snmpTable` command with the `-v1` option works as expected.

**Workaround:** Use the `-CB` option to use only GETNEXT, not GETBULK, requests to retrieve data. See “[Querying the Oracle VM Server for SPARC MIB](#)” in *Oracle VM Server for SPARC 3.1 Administration Guide*.

## SR-IOV Issues

### Bad Trap Panic Occurs Rarely When Rebooting an Oracle Solaris 10 Root Domain That Has SR-IOV Virtual Functions Assigned to Guest Domains

**Bug ID 18323562:** An Oracle Solaris 10 root domain might panic when rebooting. The Oracle Solaris 10 root domain has at least two PCIe buses and the virtual functions from the physical functions in different buses are assigned to guest domains. That is, if events from different buses on guest domains are received in parallel, the root domain might panic. This panic occurs rarely.

```
panic[cpu3]/thread=2a100365c80: BAD TRAP: type=31 rp=2a1003652b0 addr=2000
mmu_fsr=0 occurred in module "pcie" due to an illegal access to a user
address
```

**Workaround:** None.

## prtdiag Might Cause an Oracle Solaris 10 Root Domain to Panic After Destroying SR-IOV Virtual Functions

**Bug ID 18323370:** An Oracle Solaris 10 root domain might panic if you destroy virtual functions and then run the `prtdiag` command.

The `prtdiag` command might cause a panic when attempting to access virtual function device nodes that were just destroyed:

```
panic[cpu31]/thread=2a10140bc80: Fatal error has occurred in: PCIe
fabric.(0x1)(0x43)
```

And the `prtdiag` command prints messages such as the following:

```
DEV_GET failed -1 Invalid argument 4.0.2 offset 0xff
/SYS/PCI-EM4      PCIE  fibre-channel-pciex10df,e200      --
                  /pci@600/pci@1/pci@0/pci@4/fibre-channel@0,2
```

These messages occur because the `prtdiag` command attempts to access virtual function device nodes that have been destroyed. The nodes still appear in the `picl` tree, but not in the actual device tree.

**Workaround:** To avoid the panic, add the following line to the `/etc/system` file on the Oracle Solaris 10 root domain:

```
set px:pxttool_cfg_delay_usec=25000
```

Also refresh the `picl` daemon to avoid the Invalid argument messages:

```
# svcadm refresh picl
```

## Control Domain Hangs When Stopping or Starting I/O Domains

**Bug ID 18030411:** The primary domain might hang if you stop and start I/O domains frequently and in rapid succession. As a result of this behavior, the InfiniBand HCA stops responding and causes the primary domain to hang.

If you experience this problem, you might see messages on the console or in the `messages` file that are similar to the following:

```
VF3: PF has failed
Mcnnex: HW2SW_MPT command @ failed: 0000ffff
Hermon: MAD_IFC (port 01) command failed: 0000ffff
WARNING: mcxnex0: Device Error: HCR Timeout waiting for command go bit
```

**Recovery:** To avoid this problem, do not perform unnecessary stop and start operations of the I/O domains. Instead, perform an orderly shutdown of the I/O domain.

**Workaround:** If the primary domain hangs for this reason, reset the system in one of the following ways:

- Perform a reboot of the domain
 

```
primary# ldm stop -r domain-name
```
- Perform a reset in the SP
 

```
-> reset /SYS
```

## Warnings Appear on Console When Creating Fibre Channel Virtual Functions

**Bug ID 17623156:** When you create Fibre Channel virtual functions, you might see the following warnings:

```
WARNING: kmem_cache_destroy: 'px0_emlxs3_3_cache2'
(3000383e030) not empty
WARNING: vmem_destroy('px0_emlxs3_3_vmem_top'):
leaked 262144 identifiers
```

These messages do not affect the normal operation of the system and you can ignore them.

**Workaround:** None.

## Fibre Channel Physical Function Configuration Changes Require Several Minutes to Complete

**Bug ID 16397888:** After you add or destroy virtual functions, it might take up to five minutes before you can attempt to add or destroy more virtual functions from the Fibre Channel physical function.

If you attempt to perform these operations before five minutes elapse, the operations fail with a message similar to the following:

```
The attempt to offline the pf /SYS/PCI-EM4/IOVFC.PF0 in domain
primary failed.
Error message from svc:/ldoms/agents in domain primary:
CMD_OFFLINE Failed. ERROR: devices or resources are busy.
```

**Workaround:** Wait five minutes before you attempt another IOV operation on the Fibre Channel physical function.

To perform all necessary configuration options in a single command, use the `ldm create-vf -n max` or `ldm destroy-vf -n max` command.

## Fujitsu M10 System Has Different SR-IOV Feature Limitations

On a Fujitsu M10 system you can assign PCIe endpoint devices and SR-IOV virtual functions from a particular PCIe bus to a maximum of 24 domains. While on SPARC T-Series and SPARC M-Series platforms, you can assign PCIe endpoint devices and SR-IOV virtual functions from a particular PCIe bus to a maximum of 15 domains.

## InfiniBand SR-IOV Issues

---



**Caution** – Review this section before you deploy InfiniBand SR-IOV in your Oracle VM Server for SPARC 3.1 environment.

---

This section describes the known issues for the InfiniBand SR-IOV feature in the initial release of Oracle VM Server for SPARC 3.1.

- The reboot of an Oracle Solaris 11.1.10.5.0 I/O domain that has InfiniBand virtual functions assigned to it occasionally panics the corresponding root domain. See bug ID 17336355.
- An Oracle Solaris 10 1/13 I/O domain that has InfiniBand virtual functions assigned to it sometimes panics during reboot. The I/O domain runs the Oracle Solaris 10 1/13 OS plus the required patches. See bug IDs 17382933, 17361763, 17329218, and 17336035.

## Misleading Messages Shown For InfiniBand SR-IOV Operations

**Bug ID 16979993:** An attempt to use dynamic SR-IOV operations on an InfiniBand device results in confusing and inappropriate error messages.

Dynamic SR-IOV is not supported for InfiniBand devices.

**Workaround:** Manage InfiniBand virtual functions by performing the one of the following procedures:

- “How to Create an InfiniBand Virtual Function” in *Oracle VM Server for SPARC 3.1 Administration Guide*
- “How to Destroy an InfiniBand Virtual Function” in *Oracle VM Server for SPARC 3.1 Administration Guide*
- “How to Remove an InfiniBand Virtual Function From an I/O Domain” in *Oracle VM Server for SPARC 3.1 Administration Guide*
- “How to Remove an InfiniBand Virtual Function From a Root Domain” in *Oracle VM Server for SPARC 3.1 Administration Guide*

## Bugs Affecting the Oracle VM Server for SPARC Software

The following sections summarize the bugs that you might encounter when using each version of the Oracle VM Server for SPARC 3.1 software. Each section lists the bugs found in that release. The bugs might occur in any or all of the Oracle VM Server for SPARC 3.1 releases. The most recent bugs are described first. Workarounds and recovery procedures are specified, if available.

---

**Note** – Some of the bugs described in this section have been fixed since the Oracle VM Server for SPARC 3.1 release. These bug writeups remain for those who are still running the Oracle VM Server for SPARC 3.1 release.

---

### Bugs Affecting the Oracle VM Server for SPARC 3.1.1.2 Software

#### System Crashes When Applying the Whole-Core Constraint to a Partial Core primary Domain

**Bug ID 19456310:** When using dynamic reconfiguration to apply the whole-core constraint to a primary domain, the removal of partial cores results in an OS panic or system powercycle.

A partial core is removed if the core is shared with another domain or if one of the free strands in the core is faulty.

**Workaround:** Use a delayed reconfiguration to apply the whole-core constraint to a primary domain that has partial cores.

1. Verify that the primary domain does not have the whole-core constraint.

```
primary# ldm list -o resmgmt primary
```

2. Verify that the primary domain has partial cores.

```
primary# ldm list -o core primary
```

3. Initiate a delayed reconfiguration on the primary domain.

```
primary# ldm start-reconf primary
```

4. Apply the whole-core constraint.

For example, the following command assigns two whole cores to the primary domain:

```
primary# ldm set-core 2 primary
```

5. Reboot the primary domain.

## format Command Hangs After Having Migrated a Guest Domain or a Guest Domain Console Does Not Take Input

You might encounter the following bugs if your system runs system firmware version 8.5.1.b, 9.2.1.b, or 9.2.1.c. For more information, see [Oracle Virtual Machine \(OVM\) Server for SPARC Guest Domains may not Accept Console Input on SPARC T4/T5/M5/M6 Series Servers Running Sun System Firmware Releases 8.5.1.b and 9.2.1.B/C \(Doc ID 1946535.1\)](https://support.oracle.com/epmos/faces/DocumentDisplay?id=1946535.1) (<https://support.oracle.com/epmos/faces/DocumentDisplay?id=1946535.1>).

- **Bug ID 19430884:** A guest domain that is configured with 108 virtual disks from two service domains is migrated. After the migration completes successfully, the `format` command hangs even though the disks are available and can be accessed.

**Workaround:** Reboot the system.

- **Bug ID 19388985:** Attempting to connect to a guest domain console succeeds but the console does not take input. This situation occurs intermittently after stopping and starting guest domains, rebooting the primary domain, and binding and starting the guest domains.

**Workaround:** Avoid unbinding and then rebinding the guest domain.

**Recovery:** Save the configuration of the guest domains and then perform a powercycle.

## Kernel Zones Block Live Migration of Guest Domains

**Bug ID 18289196:** On a SPARC system, a running kernel zone within an Oracle VM Server for SPARC domain will block live migration of the guest domain. The following error message is displayed:

```
Live migration failed because Kernel Zones are active.
Stop Kernel Zones and retry.
```

**Workaround:** Choose one of the following workarounds:

- Stop running the kernel zone.  

```
# zoneadm -z zonename shutdown
```
- Suspend the kernel zone.  

```
# zoneadm -z zonename suspend
```

## Bugs Affecting the Oracle VM Server for SPARC 3.1.1.1 Software

### Live Migration Might Fail With Unable to restore ldc resource state on target Domain Migration of LDom failed

**Bug ID 19454837:** A live migration of a domain on a system that runs particular versions of SPARC system firmware might fail with the following error message:

```
system1 # ldm migrate ldg1 system2
Target Password:
Unable to restore ldc resource state on target
Domain Migration of LDom ldg1 failed
```

The error message occurs after transferring all the domain state to the target machine but before attempting to suspend the domain to be migrated on the source machine. The domain to be migrated continues to run on the source system.

The following are the affected system firmware versions:

- **SPARC T5, SPARC M5, SPARC M6** – System firmware version 9.2.1
- **SPARC T4** – System firmware version 8.5.1

**Mitigation:** Unless you want to take advantage of the new increased LDC limits (and not use the live migration feature), avoid updating your system to system firmware versions 8.5.1 or 9.2.1 until at least versions 8.6 and 9.3 are released.

**Recovery:** Perform a powercycle of the source machine to permit the live migration of the domain.

**Workaround:** None.

## Recovery Mode Fails With ldmd in Maintenance Mode When Virtual Switch net-dev Is Missing

**Bug ID 18770805:** If a virtual switch net - dev is faulty and cannot be validated, the recovery operation fails and the ldmd daemon dumps core.

**Recovery:** Disable recovery mode and recover the configuration manually.

## Migration to a SPARC M5 or SPARC T5 System Might Panic With suspend: get stick freq failed

**Bug ID 16934400:** When you migrate a guest domain to a SPARC M5 or a SPARC T5 system, the OS on the guest domain might panic with the suspend: get stick freq failed message.

**Workaround:** Add the following line to the /etc/system file in the guest domain to be migrated:

```
set migmd_buf_addl_size = 0x100000
```

Reboot the guest domain to make the change take effect.

## Logical Domains Manager Does Not Prohibit the Creation of Circular Dependencies

**Bug ID 15751041:** The Logical Domains Manager permits you to create a circular configuration where two domains provide services to each other. Such a configuration is not recommended because it creates a single-point-of-failure outage where one domain takes down the other domain. In addition, a circular dependency prevents the affected domains from being unbound.

**Workaround:** If a circular dependency configuration prevents you from unbinding a domain, remove the devices that cause the circular dependency and then retry the unbind operation.

## Bugs Affecting the Oracle VM Server for SPARC 3.1.1 Software

### Very Large LDC Counts Might Result in Oracle Solaris Issues in Guest Domains

**Bug ID 19480835:** The following Sun System Firmware versions increase the maximum number of Logical Domain Channels (LDCs) per guest domain:

- SPARC T5, SPARC M5, SPARC M6 – 9.2.1
- SPARC T4 – 8.5.1

This increase in LDCs per guest domain requires that you run at least Logical Domains Manager 3.1.1.1.

To prevent potential issues when using versions of the Logical Domains Manager prior to and including 3.1.1, do not increase the number of LDCs per guest domain beyond the 768 that is supported by previous system firmware versions. For example, do not add large numbers of virtual disks and virtual network interfaces until after you install at least Logical Domains Manager 3.1.1.1.

You might see the following possible symptoms when exceeding the 768 LDCs per domain limit with Oracle VM Server for SPARC versions prior to and including 3.1.1:

- Dictionary overflow in OBP:

```
Dictionary overflow - here f21ffe58 limit f2200000
```

```
Dictionary overflow - here f21ffe70 limit f2200000
```

```
WARNING: /virtual-devices@100/channel-devices@200/disk@5b2: Problem  
creating devalias for virtual device node
```

```
Dictionary overflow - here f21ffe70 limit f2200000
```

```
Dictionary overflow - here f21ffe70 limit f2200000
```

```
Dictionary overflow - here f21ffe70 limit f2200000
```

```
Stack Underflow
```

```
ok
```

- Panic in vmem\_xalloc:

```
panic[cpu6]/thread=2a10020fc80: vmem_xalloc(1a04610, 29360128, 29360128, 0,
0, 0, 0, 1): parameters inconsistent or invalid

000002a10020f000 genunix:vmem_xalloc+850 (1a04610, 1c00000, 0, 0, 1bffff, 0)
  %l0-3: 0000000000001fff 000000000002000 000000000420000 000000000000010
  %l4-7: 000000001c00000 000000000000008 000000001c00000 000000000000000
000002a10020f180 unix:contig_vmem_xalloc_aligned_wrapper+24 (1a04610,
1c00000, 1, 0, 1000000, 1)
  %l0-3: 000002a10020f9a4 000000000000008 000000001a4bd90 000000000000018
  %l4-7: 000000000000002 ffffffff0000000 00000000136efe8 0000000013722c0
000002a10020f240 genunix:vmem_xalloc+5c8 (300150c2d98, 1c00000, 0, 0, 80000,
0)
  %l0-3: 00000300150c2ff0 ffffffff0000000 00000300150c39e0 ffffffff0000000
  %l4-7: 000000000000000 ffffffff0000000 000000001000000 000000000000004
000002a10020f3c0 unix:contig_mem_span_alloc+24 (300150c2d98, 1000000, 1, 1,
cd4000, 3)
  %l0-3: 000000000000f4000 000000000000000 000000000000000 000000001921897
  %l4-7: 000000000000006 00000000fe53dce8 00000000fee3a844 000000007ffffa4c
000002a10020f490 genunix:vmem_xalloc+5c8 (300150c4000, cd4000, 0, 0, 80000,
0)
  %l0-3: 00000300150c4258 ffffffff0000000 00000300150c4c48 ffffffff0000000
  %l4-7: 000000000000000 ffffffff0000000 000000000020000 000000000000003
000002a10020f610 unix:contig_mem_alloc_align+28 (cd4000, 2000, 600957feaf8,
1, 600957feaf8, 18e3000)
  %l0-3: 0000000000000001 0000000000003000 00000300051c01d8 000000000000000
  %l4-7: 0000000000002000 000000001a29e20 00000300051c01b0 00000300051c0380
000002a10020f6d0 unix:mach_descrip_buf_alloc+8 (cd4000, 2000, 4, 1,
2a10020f838, 10448d0)
  %l0-3: 0000000000000000 0000000000003000 00000300002141d8 000000000000000
  %l4-7: 000000000000001 000000000000100 00000300002141b0 0000030000214380
000002a10020f780 unix:mach_descrip_update+84 (1864c00, 1c00, cd4000, 18e31d8,
0, 0)
  %l0-3: 000000001864c58 000002a10020f830 0000000000002000 ffffffff0000000
  %l4-7: 000002a10020f838 000000000cd27b0 000000001864c30 00000600957feaf8
000002a10020f840 platsvc:ps_md_data_handler+30 (1a4bcc0, 3003a822be0, 8, 18,
10, 1)
  %l0-3: 0000000000001d03 000000000420000 000000000420000 000000000000010
  %l4-7: 000003003a822bd8 000000000000008 000000000000008 000003000d9bb940
000002a10020f900 ds:ds_dispatch_event+30 (6009fef4df8, 1372000, 48, 9, 9,
3003a822bd0)
  %l0-3: 000002a10020f9a4 000000000000008 000000001a4bd90 000000000000018
  %l4-7: 000000000000002 ffffffff0000000 00000000136efe8 0000000013722c0
000002a10020f9b0 genunix:taskq_thread+3cc (600957fd390, 600957fd328,
260fe5123efd, 600957fd35a, 260fe5124083, 600957fd35c)
  %l0-3: 00000600957feaf8 00000600957fd358 000000000000001 000000000080000
  %l4-7: 00000600957fd348 00000000010000 00000000ffffeffff 00000600957fd350
```

## Fibre Channel Physical Function Is Faulted by FMA And Disabled

**Bug IDs 18168525 and 18156291:** You must connect the Fibre Channel PCIe card to a Fibre Channel switch that supports NPIV and that is compatible with the PCIe card. If you do not use this configuration, using the format command, or creating or destroying a virtual function might cause the physical function to be faulted by FMA and disabled. If this fault occurs, the message is similar to the following:

SUNW-MSG-ID: PCIEX-8000-0A, TYPE: Fault, VER: 1, SEVERITY: Critical  
EVENT-TIME: *event-time*  
PLATFORM: *platform-type*  
SOURCE: eft, REV: 1.16  
EVENT-ID: *event-ID*  
DESC: A problem was detected for a PCIEX device.  
AUTO\_RESPONSE: One or more device instances may be disabled  
IMPACT: Loss of services provided by the device instances associated with this fault  
REC-ACTION: Use 'fmadm faulty' to provide a more detailed view of this event. Please refer to the associated reference document at <http://support.oracle.com/msg/PCIEX-8000-0A> for the latest service procedures and policies regarding this diagnosis.

**Workaround:** If the card has been faulted by FMA, first check its connections and ensure that the card is not directly connected to storage. Then, perform the step that matches your configuration:

- **Card is directly connected to storage** – Correctly configure the Fibre Channel PCIe card by connecting it to a Fibre Channel switch that supports NPIV and is compatible with the PCIe card. Then, run the `fmadm repair` command to override the FMA diagnosis.
- **Card is not directly connected to storage** – Replace the card.

## Virtual Network LDC Handshake Issues Seen When There Are a Large Number of Virtual Network Devices Present

**Bug ID 18166010:** You might experience virtual network LDC handshake issues if your deployment has a large number of virtual network devices.

**Workaround:** Perform the following steps:

1. Increase the number of handshake retries on all domains that have a virtual network device by adding the following entry to the `/etc/system` file:

```
set vnet:vgen_ldc_max_resets = 25
```

Note that you must reboot any domain on which you updated the `/etc/system` file for the changes to take effect. For information about `/etc/system` tuneables, see the `system(4)` man page.

2. Disable inter-vnet links when a large number of virtual network devices are required in a virtual switch.

If more than eight virtual network devices use a given virtual switch, set the `inter-vnet-link` property to `off`. Disabling the `inter-vnet-link` property avoids the use of  $N^2$  channels for inter-vnet communications. This change might negatively affect the performance of inter-vnet communications. So, if the guest-to-guest performance is critical for your deployment, create a separate system-private virtual switch (without specifying a `net-dev` device) that uses only the virtual network devices that require inter-vnet communications.

If your deployment does not require high-performance guest-to-guest communications, set the `inter-vnet-link` property to `off` even if fewer virtual network devices use a given virtual switch.

```
primary# ldm set-vsw inter-vnet-link=off vsw0
```

If this workaround does not solve your problem, as a last result, make the following changes to the `/etc/system` file on all domains that have virtual network and virtual switch devices.

Note that updating the `/etc/system` file in this way might negatively affect guest-to-guest communication performance.

1. Add the following entry to the `/etc/system` file of a domain that has a virtual network device:

```
set vnet:vnet_num_descriptors = 512
```

2. Add the following entry to the `/etc/system` file of a domain that has a virtual switch device:

```
set vsw:vsw_num_descriptors = 512
```

3. Reboot the system for these settings to take effect.

## Sun Storage 16 Gb Fibre Channel Universal HBA Firmware Does Not Support Bandwidth Controls

**Bug ID 18083904:** The firmware for Sun Storage 16 Gb Fibre Channel Universal HBA, Emulex cards does not support setting bandwidth controls. The HBA firmware ignores any value that you specify for the `bw-percent` property.

**Workaround:** None.

## Adding Memory After Performing a Cross-CPU Migration Might Cause a Guest Domain Panic

**Bug ID 18032944:** Performing a cross-CPU live migration of a domain from a SPARC T5, SPARC M5, or SPARC M6 machine to a platform that runs a different CPU type succeeds. However, a subsequent memory dynamic reconfiguration operation to increase the memory size of the guest domain might cause a panic similar to the following:

```
panic[cpu0]/thread=2a1003c9c60: kphysm_add_memory_dynamic(1018000, 200000):
range has 2097152 pages, but memgr p_walk_pfnrange only reported 0
000002a1003c9500 genunix:kphysm_add_memory_dynamic+254 (1018000, 200000,
12e8000, 3, 1218000, 0)

vpanic(12e8220, 1018000, 200000, 200000, 0, 2a1003c95c8)
kphysm_add_memory_dynamic+0x254(1018000, 200000, 12e8000, 3, 1218000, 0)
dr_mem_configure+0x94(1018000, 2a1003c97b4, ffffffff, 2430000000, 1068ac00,
1068ac00)
dr_mem_list_wrk+0x15c(4c01b3382b8, 0, 20, 4c014ba27c8, 1, 1)
dr_mem_data_handler+0xa8(0, 4c01b3382b8, 20, 2a1003c9890, 7bac0644, 16)
```

```
ds_dispatch_event+0x2c(4c01ee33478, 7bf888b8, 48, 7bf88800, 9, 9)
taskq_thread+0x3a8(95af9e15e84, 4c010a5caf0, 95af9e15f74, 4c010a5cb22,
4c010a5cb24, 4c01e24d688)
thread_start+4(4c010a5caf0, 0, 0, 0, 0, 0)
```

This panic occurs when the target system is one of the following:

- SPARC T-Series systems that have socket 0 disabled
- SPARC M-Series systems that have socket 0 disabled
- Physical domains on a SPARC M-Series system that do not contain DCU0

This situation does not affect migrations between systems that have the same CPU type or domains that have `cpu-arch=native`.

**Workaround:** After migrating a domain from a system with one of these configurations, you must reboot the guest domain before you attempt to add memory by means of dynamic reconfiguration.

## Incorrect Device Path for Fibre Channel Virtual Functions in a Root Domain

**Bug ID 18001028:** In the root domain, the Oracle Solaris device path for a Fibre Channel virtual function is incorrect.

For example, the incorrect path name is `pci@380/pci@1/pci@0/pci@6/fibre-channel@0,2` while it should be `pci@380/pci@1/pci@0/pci@6/SUNW,emlxs@0,2`.

The `ldm list-io -l` output shows the correct device path for the Fibre Channel virtual functions.

**Workaround:** None.

## ldmd Dumps Core When Attempting to Bind a Domain in Either the Binding or Unbinding State

**Bug ID 17796639:** When running Oracle Enterprise Manager Ops Center 12c Release 1 Update 4 (12.1.4.0.0), if you attempt a bind, unbind, start, or stop operation on a domain that is in the binding or unbinding state, the `ldmd` service might dump core and the domain will drop to maintenance mode.

**Recovery:** If the `ldmd` service has dumped core already, perform a power cycle of the system to bring the `ldmd` service online again.

**Workaround:** Determine whether the domain is in a binding or unbinding state by running the `ldm list` command. If it is, wait until the process completes and the domain is in the bound or inactive state.

## Bugs Affecting the Oracle VM Server for SPARC 3.1 Software

### Issues Might Arise When FMA Detects Faulty Memory

**Bug IDs 17663828 and 17576087:** When FMA attempts to isolate an extremely small range of memory as a percentage of the total memory capacity of the system, the Logical Domains Manager might incorrectly mark a very large range of memory as being blacklisted.

This error can have a significant impact on usable memory capacity, which might lead to the following issues:

- The reboot of an affected guest domain might prevent that domain from starting because too much memory has been incorrectly removed.
- A very large range of memory might be unavailable for assignment to guest domains if a blacklist request is applied to unbound memory. Thus, if you attempt to use most of the system memory, you might be unable to create guest domains.
- The Logical Domains Manager might crash if it restarts prior to the faulty memory being repaired because the blacklisted memory block might not have been properly marked internally.
- Performing a power cycle of the system to a saved configuration might not restore the memory after the faulty memory has been replaced.

**Workaround:** If a significant amount of memory no longer appears in `ldm list-devices -a` memory output, contact Oracle Service to confirm and identify the DIMM that must be replaced.

After you replace the faulty memory, perform a power cycle of the system to the factory-default configuration. Then, perform a power cycle of the system to the configuration that you want to use.

### ldmd Service Fails to Start Because of a Delay in Creating virtual-channel@0:hvctl

**Bug ID 17627526:** Sometimes during a system boot, a race condition occurs where the device the `ldmd` daemon uses to communicate with the hypervisor is not created by the time the `svc:/ldoms/ldmd:default` SMF service starts. This behavior causes the `ldmd` SMF service to drop to maintenance mode.

The following error message appears in the `ldmd` SMF log:

```
ldmd cannot communicate with the hypervisor as the required device
does not exist:
/devices/virtual-devices@100/channel-devices@200/virtual-channel@0:hvctl
```

This problem might occur if the control domain is running one of the following OS versions:

- At least Oracle Solaris 11.1.12.3.0
- At least Oracle Solaris 10 1/13 and patch ID 150840-01

**Recovery:** Verify that the

/devices/virtual-devices@100/channel-devices@200/virtual-channel@0: hvctl device exists and then run the `svcadm clear ldmd` command.

## Poor Affinity on the Control Domain When You Assign Memory Before You Assign CPUs in a Delayed Reconfiguration

**Bug ID 17606070:** If you assign memory prior to assigning CPUs to the primary domain while in a delayed reconfiguration, the memory will have affinity to the allocated CPUs at the time you issue the `ldm set-memory` command even if you perform additional `ldm set-vcpu` or `ldm set-core` commands. For example, the following commands might create a situation where the 16 Gbytes of memory allocated to the primary domain might not have affinity to the eight cores that are subsequently allocated by the `ldm set-core` command:

```
primary# ldm start-reconf primary
primary# ldm set-mem 16G primary
primary# ldm set-core 8 primary
primary# reboot
```

**Workaround:** Ensure that you assign the cores to the primary domain before you assign the memory. For example, the following commands first assign eight cores to the primary domain and then assign 16 Gbytes of memory:

```
primary# ldm start-reconf primary
primary# ldm set-core 8 primary
primary# ldm set-mem 16G primary
primary# reboot
```

## Cannot Install the Oracle Solaris 11.1 OS Using an EFI GPT Disk Label on Single-Slice Virtual Disk

**Bug ID 17422973:** The installation of the Oracle Solaris 11.1 OS on a single-slice disk might fail with the following error on a SPARC T4 server that runs at least system firmware version 8.4.0, a SPARC T5, SPARC M5, or SPARC M6 server that runs at least system firmware version 9.1.0, or a Fujitsu M10 system that runs at least XCP version 2230:

```
cannot label 'cld0': try using fdisk(1M) and then provide a specific slice
Unable to build pool from specified devices: invalid vdev configuration
```

**Workaround:** Relabel the disk with an SMI label.

## After Being Migrated, A Domain Can Panic on Boot After Being Started or Rebooted

**Bug ID 17285811:** A guest domain that was previously migrated might fail to reboot on subsequent reboots or domain start operations because of a kernel panic. The panic occurs as the domain boots. The panic error message is similar to the following message:

```
panic[cpu0]/thread=10012000: tilelet_assign_cb: assigning pfns [50000, c0000)
  to mgid 1, mnodeid 1: pachunk 1 already assigned to mgid 0, mnodeid 0
```

**Workaround:** Do not reboot the domain. First, stop and unbind the domain and then bind and start the domain again. For example:

```
primary# ldm stop domain
primary# ldm unbind domain
primary# ldm bind domain
primary# ldm start domain
```

**Recovery:** When the problem occurs, stop and unbind the domain and then bind and start the domain again.

## Size of Preallocated Machine Description Buffer Is Used During Migration

**Bug ID 17285745:** Migrating a guest domain to a SPARC T5, SPARC M5, or SPARC M6 system might result in a kernel panic on the guest domain with the suspend: get stick freq failed message.

**Workaround:** Add the following setting to the /etc/system file in the guest domain that to be migrated. Then, reboot the guest domain.

```
set migmd_buf_addl_size = 0x100000
```

## Attempting to Resize a Guest Domain's Virtual CPUs After a Successful Core Remap Operation Might Fail

**Bug ID 17245915:** When FMA detects a faulty core, the Logical Domains Manager attempts to evacuate it by performing a core remap operation if a core is free to use as a target. After the core remap operation succeeds and the faulty core is replaced, attempting to resize a guest domain's virtual CPUs by using the `ldm add-vcpu` command might fail with an `Invalid response` error message.

The failure is intermittent and depends on the system configuration.

**Workaround:** None.

**Recovery:** Perform the following steps to add more CPUs to the guest domain:

1. Unbind the guest domain.

2. Remove all the virtual CPUs.
3. Add the virtual CPUs again.
4. Bind the guest domain.

The ability to reliably use DR to add CPUs will be fully restored when the blacklisted CPU resources are repaired.

## Oracle Solaris 10: Non-primary Root Domain Hangs at Boot on a primary Reboot When failure-policy=reset

**Bug ID 17232035:** A slave domain might hang on boot when `failure-policy=reset` in the master domain. This issue is not reproducible with different settings of the `failure-policy` property.

**Recovery:** Stop the I/O domains that are associated with this root domain and start the non-primary root domain.

**Workaround:** Set the `failure-policy` property to a value other than `reset`.

## Virtual Network Hang Prevents a Domain Migration

**Bug ID 17191488:** When attempting to migrate a domain from a SPARC T5-8 to a SPARC T4-4 system, the following error occurs:

```
primary# ldm migrate ldg1 system2
Target Password:
Timeout waiting for domain ldg1 to suspend
Domain Migration of LDom ldg1 failed
```

**Workaround:** To avoid this problem, set `extended-mapin-space=on`.

---

**Note** – This command initiates a delayed reconfiguration if *ldom* is primary. In all other cases, stop the domain before you perform this command.

---

```
primary# ldm set-domain extended-mapin-space=on ldom
```

## ldmpower Output Sometimes Does Not Include Timestamps

**Bug ID 17188920:** The `--suppress` and `--timestamp` options do not properly show timestamp values.

**Workaround:** Include the `-r` option when using the `--suppress` and `--timestamp` options to show the correct output.

## mac\_do\_softlso Drops LSO Packets

**Bug ID 17182503:** `mac_do_softlso()` drops LSO packets that are generated by the `vnet_vlan_insert_tag()` and `vnet_vlan_remove_tag()` functions.

**Workaround:** To avoid this issue with VLAN tagged LSO packets, disable virtual network LSO capability on all domains that support it.

1. Append the following lines to the `/etc/system` file:

```
set vnet_enable_lso = 0
set vsw_enable_lso = 0
```

2. Reboot.
3. Verify the changes by using the `mdb -k` command.

```
# mdb -k
> vnet_enable_lso/D
vnet_enable_lso:
vnet_enable_lso:0

> vsw_enable_lso/D
vsw_enable_lso:
vsw_enable_lso: 0
```

## Migration Failure: Invalid Shutdown-group: 0

**Bug ID 17088083:** The migration of a domain that has more than eight virtual CPUs might result in memory corruption if the domain's highest processor group ID increases across a 64-unit multiple. For example, before the migration the highest processor group ID on the domain is 63 and after the migration it is 64.

Use the `pginfo` command to determine the processor group IDs in a domain. Within a domain, run the following command to print the highest processor group ID:

```
# pginfo -I|tr ' ' '\n'|sort -n|tail -1
```

**Workaround:** Reduce the number of virtual CPUs in the domain to eight before performing the migration. After the migration completes, you can restore the virtual CPU count in the domain to the original value.

## Autosave Configuration Is Not Updated After the Removal of a Virtual Function or a PCIe Device

**Bug ID 17051532:** When a PCIe device or a virtual function is removed from a guest domain, the autosave configuration is not updated. This problem might result in the device or virtual function reappearing in the guest domain after you perform an autosave recovery; namely, when `autorecovery_policy=3`. This problem can also cause the `ldm add-sponconfig -r` command to fail with the Autosave configuration `config-name` is invalid message if you do not perform another `ldm` command that causes the autosave to be updated.

**Workaround:** Perform one of the following workarounds:

- Save a new configuration after you remove the PCIe device or virtual function.

```
primary# ldm add-config new-config-name
```

- Refresh the saved configuration after you remove the PCIe device or virtual function by removing and re-creating the configuration.

```
primary# ldm rm-config config-name
primary# ldm add-config config-name
```

Note that this bug prevents the `ldm add-config -r config-name` command from working properly.

- Issue another `ldm` command that causes an autosave update to occur such as `ldm set -vcpu`, `ldm bind`, or `ldm unbind`.

### ldmp2v convert Command Failure Causes Upgrade Loop

**Bug ID 17026219:** If an error occurs during the `ldmp2v convert` command, sometimes the `boot-device` property for the guest is not set to the guest's boot disk. This error causes the guest domain to boot from the Oracle Solaris Install image again after the Oracle Solaris upgrade finishes.

**Workaround:** Change the `boot-device` property on the guest domain from within the control domain. Make this change when you re-enter the Oracle Solaris Installer, and then redo the Oracle Solaris upgrade. The guest domain will then reboot from the upgraded boot disk after the upgrade has finished.

To set the boot device, run the following command on the control domain. This command assumes that the root (`/`) file system of the original physical system is located on slice 0 of the boot disk. If the original system booted from another slice, adjust the letter after the colon accordingly. For instance, use `a` for slice 0, `b` for slice 1, and so on.

```
primary# ldm set-variable boot-device=disk0:a domain-name
```

### Domain Migrations From SPARC T4 Systems That Run System Firmware 8.3 to SPARC T5, SPARC M5, or SPARC M6 Systems Are Erroneously Permitted

**Bug ID 17027275:** Domain migrations between SPARC T4 systems that run system firmware 8.3 should not be permitted to SPARC T5, SPARC M5, or SPARC M6 systems. Although the migration succeeds, a subsequent memory DR operation causes a panic.

**Workaround:** Update the system firmware on the SPARC T4 system to version 8.4. See the workaround for “[Guest Domain Panics at `lgrp\_lineage\_add\(mutex\_enter: bad mutex, lp=10351178\)`”](#) on page 63.

## Guest Domain Panics at `lgrp_lineage_add(mutex_enter: bad mutex, lp=10351178)`

**Bug ID 17020950:** After migrating an active domain from a SPARC T4 platform to a SPARC T5, SPARC M5, or SPARC M6 platform that was bound using firmware version 8.3, performing a memory dynamic reconfiguration might result in a guest domain panic.

**Workaround:** Before you perform the migration, update the SPARC T4 system with version 8.4 of the system firmware. Then, rebind the domain.

## Guest Domains in Transition State After Reboot of the primary Domain

**Bug ID 17020481:** A guest domain is in transition state (`t`) after a reboot of the primary domain. This problem arises when a large number of virtual functions are configured on the system.

**Workaround:** To avoid this problem, retry the OBP disk boot command several times to avoid a boot from the network.

Perform the following steps on each domain:

1. Access the console of the domain.

```
primary# telnet localhost domain-name
```

2. Set the boot-device property.

```
ok> setenv boot-device disk disk disk disk disk disk disk disk disk net
```

The number of `disk` entries that you specify as the value of the `boot-device` property depends on the number of virtual functions that are configured on the system. On smaller systems, you might be able to include fewer instances of `disk` in the property value.

3. Verify that the boot-device property is set correctly by using the `printenv`.

```
ok> printenv
```

4. Return to the primary domain console.
5. Repeat Steps 1-4 for each domain on the system.
6. Reboot the primary domain.

```
primary# shutdown -i6 -g0 -y
```

## Panic Occurs in Rare Circumstances When the Virtual Network Device Driver Operates in TxDring Mode

**Bug ID 16991255:** A panic occurs in rare circumstances when the virtual network device driver operates in TxDring mode.

**Workaround:** To avoid this panic, set the `extended-mapin-space` property value to `on`.

**Note** – This command initiates a delayed reconfiguration if *ldom* is primary. In all other cases, stop the domain before you perform this command.

---

```
primary# ldm set-domain extended-mapin-space=on ldom
```

## A Domain That Has Only One Virtual CPU Assigned Might Panic During a Live Migration

**Bug ID 16895816:** Performing the migration of a domain that has only one virtual CPU assigned to it might panic on the guest domain in the `pg_cmt_cpu_fini()` function.

**Workaround:** Assign at least two virtual CPUs to the guest domain before you migrate it. For example, use the `ldm add-vcpu 2 domain-name` command to increase the number of virtual CPUs that are assigned to the *domain-name* guest domain.

## ldm migrate -n Should Fail When Cross-CPU Migration From SPARC T5, SPARC M5, or SPARC M6 System to UltraSPARC T2 or SPARC T3 System

**Bug ID 16864417:** The `ldm migrate -n` command does not report failure when attempting to migrate between a SPARC T5, SPARC M5, or SPARC M6 machine and an UltraSPARC T2 or SPARC T3 machine.

**Workaround:** None.

## Recovery Mode Should Support PCIe Slot Removal in Non-primary Root Domains

**Bug ID 16713362:** PCIe slots cannot currently be removed from non-primary root domains during the recovery operation. The PCIe slots remain assigned to the non-primary root domain.

**Workaround:** The PCIe slots must be removed manually from the non-primary root domain and assigned to the appropriate I/O domain or domains after the recovery operation has finished.

For information about how to remove PCIe slots from a non-primary root domain, see [“Using Non-primary Root Domains”](#) in *Oracle VM Server for SPARC 3.1 Administration Guide*.

Recovering I/O domains that use PCIe slots owned by non-primary root domains depends on the I/O domain configuration:

- If the I/O domain uses only PCIe slots and none of its PCIe slots are available, the I/O domain is not recovered and is left in the unbound state with the PCIe slots marked as evacuated.

- If the I/O domain uses SR-IOV virtual functions and PCIe slots, the domain is recovered with the unavailable PCIe slots marked as evacuated.

Use the `ldm add -io` command to add the PCIe slots to an I/O domain after you have manually removed them from the non-primary root domain.

### ldm list Does Not Show the evacuated Property for Physical I/O Devices

**Bug ID 16617981:** `ldm list` output does not show the evacuated property for the physical I/O devices.

**Workaround:** Use the `-p` option with any of the `ldm list` commands to show the evacuated property for physical I/O devices.

### Invalid Physical Address Is Received During a Domain Migration

**Bug ID 16494899:** In rare circumstances, a domain migration is rejected with the following message in the `ldmd` SMF log:

```
Mar 08 17:42:12 warning: Received invalid physical address during
migration of domain rztcrmdev2: base RA: 0x400000000, offset: 0x1ffff0000,
PA: 0x87fff0000 size: 0x1001a
```

Because the migration fails before the domain is suspended on the source system, there is no loss of service.

This failure occurs when the following circumstances occur to cause the rejection of the migration:

- The memory contents of the last chunk of memory in the domain are compressed to a size larger than the memory chunk
- The `ldmd` daemon incorrectly determines that data was written to memory outside of the domain on the target

The failure mode depends on the domain workload and the exact memory contents as most chunks are compressed to a smaller size.

**Recovery:** Although no workaround is guaranteed for this problem, performing a subsequent migration might work if the workload changes and therefore the memory contents change. You might also attempt to use dynamic reconfiguration to modify the memory size of the domain.

### send\_mondo\_set: timeout Panic Occurs When Using the ldm stop Command on a Guest Domain After Stress

**Bug ID 16486383:** This problem can occur if you assign a PCI device or bus directly to a guest domain where the domain does not have a core assigned from the `/SYS/DCU` where the PCI card physically resides. Because the hypervisor resets PCI devices on behalf of guest domains, during

each guest domain reboot a domain with cores on the DCU connected to the PCI device might panic. More PCI devices assigned to non-DCU-local guests increases the possibility of panics.

**Workaround:** Perform one of the following workarounds:

- Ensure that when you assign PCI devices to a guest domain, the card is located physically in the same DCU as the cores.
- Manually assign cores for physical card placement flexibility.  
As an example, for a PCI device on IOU0 (pci\_0 through pci\_15), choose a core between 0 and 127, and allocate it to the domain.

```
# ldm add-core cid=16 domain
```

View the system cores by using the following command:

```
# ldm ls-devices -a core
```

For a PCI device on IOU1 (pci\_16 through pci\_31), choose a core between 128 and 255. For a PCI device on IOU2 (pci\_32 through pci\_47), choose a core between 256 and 383. For a PCI device on IOU3 (pci\_48 through pci\_63), choose a core between 384 and 511.

## Subdevices Under a PCIe Device Revert to an Unassigned Name

**Bug ID 16299053:** After disabling a PCIe device, you might experience unexpected behavior. The subdevices that are under the disabled PCIe device revert to the non-assigned names while the PCIe device is still owned by the domain.

**Workaround:** If you decide to disable a PCIe slot on the ILOM, ensure that the PCIe slot is not assigned to a domain by means of the direct I/O (DIO) feature. That is, first ensure that the PCIe slot is assigned to the corresponding root domain before disabling the slot on the ILOM.

If you disable the PCIe slot on the ILOM while the PCIe slot is assigned to a domain with DIO, stop that domain and reassign the device to the root domain for the correct behavior.

## WARNING: ddi\_intr\_alloc: cannot fit into interrupt pool Means That Interrupt Supply Is Exhausted While Attaching I/O Device Drivers

**Bug ID 16284767:** This warning on the Oracle Solaris console means the interrupt supply was exhausted while attaching I/O device drivers:

```
WARNING: ddi_intr_alloc: cannot fit into interrupt pool
```

The hardware provides a finite number of interrupts, so Oracle Solaris limits how many each device can use. A default limit is designed to match the needs of typical system configurations, however this limit may need adjustment for certain system configurations.

Specifically, the limit may need adjustment if the system is partitioned into multiple logical domains and if too many I/O devices are assigned to any guest domain. Oracle VM Server for

SPARC divides the total interrupts into smaller sets given to guest domains. If too many I/O devices are assigned to a guest domain, its supply might be too small to give each device the default limit of interrupts. Thus, it exhausts its supply before it completely attaches all the drivers.

Some drivers provide an optional callback routine which allows Oracle Solaris to automatically adjust their interrupts. The default limit does not apply to these drivers.

**Workaround:** Use the `::irmools` and `::irmreqs` MDB macros to determine how interrupts are used. The `::irmools` macro shows the overall supply of interrupts divided into pools. The `::irmreqs` macro shows which devices are mapped to each pool. For each device, `::irmreqs` shows whether the default limit is enforced by an optional callback routine, how many interrupts each driver requested, and how many interrupts the driver is given.

The macros do not show information about drivers that failed to attach. However, the information that is shown helps calculate the extent to which you can adjust the default limit. Any device that uses more than one interrupt without providing a callback routine can be forced to use fewer interrupts by adjusting the default limit. Reducing the default limit below the amount that is used by such a device results in freeing of interrupts for use by other devices.

To adjust the default limit, set the `ddi_msix_alloc_limit` property to a value from 1 to 8 in the `/etc/system` file. Then, reboot the system for the change to take effect.

To maximize performance, start by assigning larger values and decrease the values in small increments until the system boots successfully without any warnings. Use the `::irmools` and `::irmreqs` macros to measure the adjustment's impact on all attached drivers.

For example, suppose the following warnings are issued while booting the Oracle Solaris OS in a guest domain:

```
WARNING: emlxs3: interrupt pool too full.
WARNING: ddi_intr_alloc: cannot fit into interrupt pool
```

The `::irmools` and `::irmreqs` macros show the following information:

```
# echo "::irmools" | mdb -k
ADDR          OWNER      TYPE      SIZE  REQUESTED  RESERVED
00000400016be970 px#0      MSI-X    36    36         36

# echo "00000400016be970::irmreqs" | mdb -k
ADDR          OWNER      TYPE      CALLBACK  NINTRS  NREQ  NAVAIL
00001000143acaa8 emlxs#0  MSI-X    No        32      8     8
00001000170199f8 emlxs#1  MSI-X    No        32      8     8
000010001400ca28 emlxs#2  MSI-X    No        32      8     8
0000100016151328 igb#3    MSI-X    No        10      3     3
0000100019549d30 igb#2    MSI-X    No        10      3     3
0000040000e0f878 igb#1    MSI-X    No        10      3     3
000010001955a5c8 igb#0    MSI-X    No        10      3     3
```

The default limit in this example is eight interrupts per device, which is not enough interrupts to accommodate the attachment of the final emLxs3 device to the system. Assuming that all emLxs instances behave in the same way, emLxs3 probably requested 8 interrupts.

By subtracting the 12 interrupts used by all of the igb devices from the total pool size of 36 interrupts, 24 interrupts are available for the emLxs devices. Dividing the 24 interrupts by 4 suggests that 6 interrupts per device would enable all emLxs devices to attach with equal performance. So, the following adjustment is added to the `/etc/system` file:

```
set ddi_msix_alloc_limit = 6
```

When the system successfully boots without warnings, the `::irm pools` and `::irm reqs` macros show the following updated information:

```
# echo "::irm pools" | mdb -k
ADDR          OWNER      TYPE      SIZE  REQUESTED  RESERVED
00000400018ca868 px#0      MSI-X    36    36         36

# echo "00000400018ca868::irm reqs" | mdb -k
ADDR          OWNER      TYPE      CALLBACK  NINTRS  NREQ  NAVAIL
0000100016143218 emLxs#0   MSI-X    No        32      8     6
0000100014269920 emLxs#1   MSI-X    No        32      8     6
000010001540be30 emLxs#2   MSI-X    No        32      8     6
00001000140cbe10 emLxs#3   MSI-X    No        32      8     6
00001000141210c0 igb#3     MSI-X    No        10      3     3
0000100017549d38 igb#2     MSI-X    No        10      3     3
0000040001ceac40 igb#1     MSI-X    No        10      3     3
000010001acc3480 igb#0     MSI-X    No        10      3     3
```

## SPARC M5-32 and SPARC M6-32: panic: mpo\_cpu\_add: Cannot read MD

**Bug ID 16238762:** On a SPARC M5-32 or a SPARC M6-32 with at least 2.4 Tbytes of memory, attempting to set the number of CPUs in the primary domain from 6 to 1056 CPUs causes the kernel to panic with the following message:

```
mpo_cpu_add: Cannot read MD
```

The following procedure causes the panic:

1. Power on with a DCU assigned to a host.  
For example, assign DCU0 to HOST0.
2. Create guest domains.
3. Save a configuration to the SP.
4. Power off the host.
5. Assign another DCU to the host.  
For example, assign DCU1 to HOST0.
6. Power on the host.

The firmware verifies that the configuration is “bootable.” This verification ensures that all the CPUs, memory, and I/O that were present at the time the configuration was created are still present. The firmware also generates a new PRI to describe the configuration of the entire system.

The configuration successfully powers on and guest domains are booted.

7. Attempt to dynamically add a CPU to an existing domain.

A new machine description is generated that reflects correct latency information, but the Oracle Solaris OS cannot parse the new information and panics.

**Workaround:** To avoid the panic, do not perform the steps that are in the problem description.

If you have already performed these steps and experienced the panic, perform the following steps:

1. Perform an action after booting a saved configuration from a smaller physical domain. For example, remove a CPU from each active domain.
2. Reboot the domain.
3. Unbind the domain.
4. Rebind any bound domains.
5. Save a new configuration to the SP.

## SPARC M5-32 and SPARC M6-32: Issue With Disks That Are Accessible Through Multiple Direct I/O Paths

**Bug ID 16232834:** When using the `ldm add -vcpu` command to assign CPUs to a domain, the Oracle Solaris OS might panic with the following message:

```
panic[cpu16]/thread=c4012102c860: mpo_cpu_add: Cannot read MD
```

This panic occurs if the following conditions exist:

- Additional DCUs have been assigned to a host
- The host is started by using a previously saved SP configuration that does not contain all the hardware that is assigned to the host

The target domain of the `ldm add -vcpu` operation is the domain that panics. The domain recovers with the additional CPUs when it reboots.

**Workaround:** Do not use configurations that are generated with fewer hardware resources than are assigned to the host.

To avoid the problem, do not add CPUs as described in the problem description. Or, perform the following steps:

1. Generate a new SP configuration after the DCUs have been added.

For example, the following command creates a configuration called `new-config-more-dcus`:

```
primary# ldm add-config new-config-more-dcus
```

2. Shutdown the domain.

3. Stop the host.

```
-> stop /HOST
```

4. Start the host.

```
-> start /HOST
```

### **ixgbev Device in SR-IOV Domains Might Become Disabled When Rebooting the primary Domain**

**Bug ID 16224353:** After rebooting the primary domain, `ixgbev` instances in primary domain might not work.

**Workaround:** None.

### **Reboot of the Oracle Solaris 10 1/13 primary Domain Might Not Automatically Plumb or Assign an IP Address to a Virtual Function Interface**

**Bug ID 16219069:** On a primary domain that runs the Oracle Solaris 10 1/13 OS, the virtual function interfaces might not be automatically plumbed or assigned an IP address based on the `/etc/hostname.vf-interface` file.

This issue occurs when you boot or reboot a SPARC T3, SPARC T4 or SPARC T5 system that runs the Oracle Solaris 10 1/13 OS on the primary domain. This problem affects virtual functions that have been created both on on-board physical functions and on add-in physical functions. This issue does not occur when you boot a Logical Domains guest domain image.

### **Oracle Solaris 10 Only: mutex\_enter: bad mutex Panic in primary Domain During a Reboot or Shutdown**

**Bug ID 16080855:** During a reboot or shutdown of the primary domain, the primary domain might experience a kernel panic with a panic message similar to the following:

```
panic[cpu2]/thread=c40043b818a0: mutex_enter: bad mutex, lp=c4005fa01c88
owner=c4005f70aa80 thread=c40043b818a0
```

```
000002a1075c3630 ldc:ldc_mem_rdwr_cookie+20 (c4005fa01c80,
c4004e2c2000,2a1075c37c8, 6c80000, 1, 0)
%l0-3: 00000000001356a4 0000000000136800 0000000000000380
00000000000002ff
%l4-7: 00000000001ad3f8 0000000000000004 00000000ffbfb9c
0000c4005fa01c88
000002a1075c3710 vlvc:i_vlvc_ioctl_write_cookie+a4 (c4004c400030,
```

```

380,ffbff898, 100003, 0, 70233400)
%l0-3: 0000000006c80000 000000000156dc8 000000000000380
000000000100003
%l4-7: 00000000702337b0 000002a1075c37c8 0000000000040000
000000000000000
000002a1075c37f0 vldc:vldc_ioctl+1a4 (3101, c4004c400030,
ffbff898,c4004c400000, c4004c438030, 0)
%l0-3: 0000000000100003 0000000000000000 000000007b340400
0000c4004c438030
%l4-7: 0000c4004c400030 0000000000000000 0000000000000000
0000000000000000
000002a1075c38a0 genunix:fop_ioctl+d0 (c4004d327800, 0, ffbff898,
100003,c4004384f718, 2a1075c3acc)
%l0-3: 0000000000003103 000000000100003 000000000133ce94
0000c4002352a480
%l4-7: 0000000000000000 0000000000000002 00000000000000c0
0000000000000000
000002a1075c3970 genunix:ioctl+16c (3, 3103, ffbff898, 3, 134d50, 0)
%l0-3: 0000c40040e00a50 000000000000c6d3 0000000000000003
0000030000002000
%l4-7: 0000000000000003 0000000000000004 0000000000000000
0000000000000000

```

**Recovery:** Allow the primary domain to reboot. If the primary domain is configured not to reboot after a crash, manually boot the primary domain.

## SPARC M5-32 and SPARC M6-32: LSI-SAS Controller Is Incorrectly Exported With SR-IOV

**Bug ID 16071170:** On a SPARC M5-32 or a SPARC M6-32 system, the internal SAS controllers are exported as SR-IOV-enabled controllers even though these cards do not support SR-IOV.

The Oracle VM Server for SPARC log shows the following messages when attempting to create the physical function on these cards:

```

Dec 11 04:27:54 warning: Dropping pf
pci@d00/pci@1/pci@0/pci@0/pci@0/pci@4/LSI,sas@0: no IOV capable driver
Dec 11 04:27:54 warning: Dropping pf
pci@d80/pci@1/pci@0/pci@c/pci@0/pci@4/LSI,sas@0: no IOV capable driver
Dec 11 04:27:54 warning: Dropping pf
pci@c00/pci@1/pci@0/pci@c/pci@0/pci@4/LSI,sas@0: no IOV capable driver
Dec 11 04:27:54 warning: Dropping pf
pci@e00/pci@1/pci@0/pci@0/pci@0/pci@4/LSI,sas@0: no IOV capable driver

```

The system has four LSI SAS controller ports, each in one IOU of the SPARC M5-32 and SPARC M6-32 assembly. This error is reported for each port.

**Workaround:** You can ignore these messages. These messages indicate only that the LSI-SAS controller devices on the system are capable of SR-IOV but no SR-IOV support is available for this hardware.

## SPARC T5-8: Uptime Data Shows a Value of 0 for Some ldm List Commands

**Bug ID 16068376:** On a T5-8 with approximately 128 domains, some ldm commands such as `ldm list` might show 0 seconds as the uptime for all domains.

**Workaround:** Log in to the domain and use the `uptime` command to determine the domain's uptime.

## Cannot Set a Jumbo MTU for sxge Virtual Functions in the primary Domain of a SPARC T5-1B System

**Bug ID 16059331:** The `sxge` driver cannot properly set jumbo MTUs for its virtual functions on the primary domain.

**Workaround:** Manually modify the `/kernel/drv/sxge.conf` file to set up the jumbo MTU on `sxge` virtual function interfaces in the guest domain.

## ldmd Is Unable to Set the mac-addr and alt-mac-addrs Property Values for the sxge Device

**Bug ID 15974640:** The `ldm` command fails to set the `mac-addr` and `alt-mac-addrs` property values properly for the `sxge` device. As a result, the `ldmd` daemon reports an inconsistent MAC address. Also, any link aggregations that are based on the VNIC MAC address also fail.

## ldm list-io -d Output for an sxge Device on SPARC T5-1B System Is Missing Two Properties

**Bug ID 15974547:** When run on a SPARC T5-1B system that has `sxge`, the `ldm list-io -d PF-device` output does not show the `max-vlans` or `max-vf-mtu` properties. These properties are present on a SPARC T5-1B system with `ixgbe` as well as on non-blade systems.

The `max-vlans` property value is missing. The value should be 0 because `sxge` device does not support hardware VLAN tagging. The `max-vf-mtu` property value is fixed at 1500, which prevents the physical function driver to set the jumbo MTU for virtual functions.

## ldm Fails to Evacuate a Faulty Core From a Guest Domain

**Bug ID 15962837:** A core evacuation does not complete when a chip-level fault occurs. An evacuation that is followed by a core fault works as expected, but the chip-level fault does not complete when trying to retire an entire CMP node.

**Workaround:** None. Schedule a chip replacement when you diagnose a chip-level fault.

## Memory DR Operations Hang When Reducing Memory Below Four Gbytes

**Bug ID 15942036:** If you perform a memory DR operation to reduce memory below four Gbytes, the operation might hang forever. If you issue an `ldm cancel -op memdr` command on that domain, an incorrect message is issued:

```
The memory removal operation has completed. You cannot cancel this operation.
```

Despite the message, the memory DR operation is hung and you might not be able to perform other `ldmd` operations on that guest domain.

**Workaround:** Do not attempt to reduce memory in any domain below four Gbytes. If you are already in this state, issue the `ldm stop -f` command or log in to the domain and reboot it.

## CPU DR of Very Large Number of Virtual CPUs Can Appear to Fail

**Bug ID 15826354:** CPU dynamic reconfiguration (DR) of a very large number of CPUs causes the `ldmd` daemon to return a failure. Although `ldmd` times out, the DR operation continues in the background and eventually succeeds. Nevertheless, `ldmd` is no longer aligned with the resulting domain and subsequent DR operations might not be permitted.

For example:

```
# ldm ls
NAME          STATE      FLAGS    CONS   VCPU  MEMORY  UTIL  NORM  UPTIME
primary      active    -n-cv-  UART   7     20G     2.7%  0.4%  1h 41m
ldg0         active    -n----  5000   761   16G     75%   51%   6m
```

```
# ldm rm-vcpu 760 ldg0
```

```
Request to remove cpu(s) sent, but no valid response received
VCPU(s) will remain allocated to the domain, but might
not be available to the guest OS
Resource removal failed
```

```
# ldm set-vcpu 1 ldg0
```

```
Busy executing earlier command; please try again later.
Unable to remove the requested VCPUs from domain ldg0
Resource modification failed
```

```
# ldm ls
NAME          STATE      FLAGS    CONS   VCPU  MEMORY  UTIL  NORM  UPTIME
primary      active    -n-cv-  UART   7     20G     0.9%  0.1%  1h 45m
ldg0         active    -n----  5000   761   16G    100%  0.0%  10m
```

**Workaround:** Wait a few minutes and then run the `ldm set-vcpu` command again:

```
# ldm set-vcpu 1 ldg0
```

```
# ldm ls
NAME          STATE      FLAGS    CONS   VCPU  MEMORY  UTIL  NORM  UPTIME
primary      active    -n-cv-  UART   7     20G     0.9%  0.1%  1h 50m
ldg0         active    -n----  5000   1     16G     52%   0.0%  15m
```

Note that 760 exceeds the recommended maximum.

## Migration of a Guest Domain With HIO Virtual Networks and `cpu-arch=generic` Times Out While Waiting for the Domain to Suspend

**Bug ID 15825538:** On a logical domain that is configured with both Hybrid network I/O interfaces (`mode=hybrid`) and cross-CPU migration enabled (`cpu-arch=generic`), if a secure live migration is executed (`ldm migrate`), the migration might time out and leave the domain in a suspended state.

**Recovery:** Restart the logical domain.

**Workaround:** Do not use hybrid I/O virtual network devices with secure cross-CPU live migration.

## SPARC T4-4: Unable to Bind a Guest Domain

**Bug ID 15825330:** Oracle VM Server for SPARC appears to hang at startup on some SPARC T4-4 configurations that have only a single processor board.

**Workaround:** Ensure that a processor board always occupies the slots for processors 0 and 1. Restarting the system in such a configuration enables the Oracle VM Server for SPARC software to start up.

## Guest Domain Panics While Changing the threading Property Value From `max-throughput` to `max-ipc`

**Bug ID 15821246:** On a system that runs the Oracle Solaris 11.1 OS, changing the threading property value on a migrated domain from `max-ipc` to `max-throughput` can lead to a panic on the guest domain.

**Workaround:** Do not change the threading status for a migrated guest domain until it is rebooted.

## Control Domain Hangs on Reboot With Two Active Direct I/O Domains

**Bug ID 15820741:** On an Oracle Solaris 11.1 system that has two domains with direct I/O configurations, the control domain might hang when you reboot it.

**Recovery:** To recover from the reboot hang, reset the control domain by issuing the following command on the SP:

```
-> reset -f /HOST/domain/control
```

## No Error Message When a Memory DR Add is Partially Successful

**Bug ID 15812823:** In low free-memory situations, not all memory blocks can be used as part of a memory DR operation due to size. However, these memory blocks are included in the amount of free memory. This situation might lead to a smaller amount of memory being added to the domain than expected. No error message is shown if this situation occurs.

**Workaround:** None.

## Primary or Guest Domain Panics When Unbinding or Migrating a Guest Domain That Has Hybrid I/O Network Devices

**Bug ID 15803617:** The primary domain or an active guest domain might panic during an unbind operation or a live migration operation if the domain is configured with hybrid I/O virtual network devices.

**Recovery:** Restart the affected domain.

**Workaround:** Do not use hybrid I/O virtual network devices.

## Re-creating a Domain That Has PCIe Virtual Functions From an XML File Fails

**Bug ID 15783851:** You might encounter a problem when attempting to re-create a configuration from an XML file that incorrectly represents virtual function constraints.

This problem occurs when you use the `ldm list-constraints -x` command to save the configuration of a domain that has PCIe virtual functions.

If you later re-create the domain by using the `ldm add-domain -i` command, the original virtual functions do not exist, and a domain bind attempt fails with the following error message:

```
No free matching PCIe device...
```

Even if you create the missing virtual functions, another domain bind attempt fails with the same error message because the virtual functions are miscategorized as PCIe devices by the `ldm add-domain` command.

**Workaround:** Perform the following steps:

1. Save the information about the virtual functions by using the `ldm list-io` command.
2. Destroy each affected domain by using the `ldm rm-dom` command.
3. Create all the required virtual functions by using the `ldm create-vf` command.
4. Rebuild the domains by using the `ldm` command.

When you use the `ldm add-io` command to add each virtual function, it is correctly categorized as a virtual function device, so the domain can be bound.

For information about rebuilding a domain configuration that uses virtual functions, see [“`ldm init-system` Command Might Not Correctly Restore a Domain Configuration on Which Physical I/O Changes Have Been Made” on page 76](#).

## Incorrect Error Message Issued When Changing the Control Domain From Using Whole Cores to Using Partial Cores

**Bug ID 15783608:** When you change the control domain from using physically constrained cores to using unconstrained CPU resources, you might see the following extraneous message:

Whole-core partitioning has been removed from domain primary, because dynamic reconfiguration has failed and the domain is now configured with a partial CPU core.

**Workaround:** You can ignore this message.

## ldm init-system Command Might Not Correctly Restore a Domain Configuration on Which Physical I/O Changes Have Been Made

**Bug ID 15783031:** You might experience problems when you use the `ldm init-system` command to restore a domain configuration that has used direct I/O or SR-IOV operations.

A problem arises if one or more of the following operations have been performed on the configuration to be restored:

- A slot has been removed from a bus that is still owned by the primary domain.
- A virtual function has been created from a physical function that is owned by the primary domain.
- A virtual function has been assigned to the primary domain, to other guest domains, or to both.
- A root complex has been removed from the primary domain and assigned to a guest domain, and that root complex is used as the basis for further I/O virtualization operations.

In other words, you created a non-primary root domain and performed any of the previous operations.

To ensure that the system remains in a state in which none of the previous actions have taken place, see [Using the ldm init-system Command to Restore Domains on Which Physical I/O Changes Have Been Made \(https://support.oracle.com/epmos/faces/DocumentDisplay?id=1575852.1\)](https://support.oracle.com/epmos/faces/DocumentDisplay?id=1575852.1).

## Logical Domains Manager Might Crash and Restart When You Attempt to Modify Many Domains Simultaneously

**Bug ID 15782994:** Logical Domains Manager might crash and restart when you attempt an operation that affects the configuration of many domains. You might see this issue when you attempt to change anything related to the virtual networking configuration and if many virtual network devices in the same virtual switch exist across many domains. Typically, this issue is seen with approximately 90 or more domains that have virtual network devices connected to

the same virtual switch, and the `inter-vnet-link` property is enabled (the default behavior). Confirm the symptom by finding the following message in the `ldm` log file and a core file in the `/var/opt/SUNWldm` directory:

```
Frag alloc for 'domain-name'/MD memory of size 0x80000 failed
```

**Workaround:** Avoid creating many virtual network devices connected to the same virtual switch. If you intend to do so, set the `inter-vnet-link` property to `off` on the virtual switch. Be aware that this option might negatively affect network performance between guest domains.

## ldm list -o Command No Longer Accepts Format Abbreviations

**Bug ID 15781142:** The `ldm list -o format` command no longer accepts abbreviations for *format*.

Although the Oracle VM Server for SPARC 3.0 software enabled you to use the `ldm list -o net` command to show information about the network, such abbreviations have been removed from the Oracle VM Server for SPARC 3.1 software. In Oracle VM Server for SPARC 3.1, you must use the full version of *format* in the command: `ldm list -o network`.

**Workaround:** Use the format names that are specified in the [ldm\(1M\)](#) man page.

## Control Domain Requires the Lowest Core in the System

**Bug ID 15778392:** The control domain requires the lowest core in the system. So, if core ID 0 is the lowest core, it cannot be shared with any other domain if you want to apply the whole-core constraint to the control domain.

For example, if the lowest core in the system is core ID 0, the control domain should look similar to the following output:

```
# ldm ls -o cpu primary
NAME
primary

VCPUs
VID  PID  CID  UTIL  STRAND
0    0    0    0.4%  100%
1    1    0    0.2%  100%
2    2    0    0.1%  100%
3    3    0    0.2%  100%
4    4    0    0.3%  100%
5    5    0    0.2%  100%
6    6    0    0.1%  100%
7    7    0    0.1%  100%
```

## After Canceling a Migration, ldm Commands That Are Run on the Target System Are Unresponsive

**Bug ID 15776752:** If you cancel a live migration, the memory contents of the domain instance that is created on the target must be “scrubbed” by the hypervisor. This scrubbing process is performed for security reasons and must be complete before the memory can be returned to the pool of free memory. While this scrubbing is in progress, ldm commands become unresponsive. As a result, the Logical Domains Manager appears to be hung.

**Recovery:** You must wait for this scrubbing request to finish before you attempt to run other ldm commands. This process might take a long time. For example, a guest domain that has 500 Gbytes of memory might complete this process in up to 7 minutes on a SPARC T4 server or up to 25 minutes on a SPARC T3 server.

## Some Emulex Cards Do Not Work When Assigned to I/O Domain

**Bug ID 15776319:** On a system that runs the Oracle Solaris OS on the control domain and an I/O domain, some Emulex cards that are assigned to the I/O domain do not function properly because the cards do not receive interrupts. However, when assigned to the control domain, the same cards work properly.

This problem occurs with the following Emulex cards:

- Emulex 2-Gigabit/Sec PCI Express Single and Dual FC Host Adapter (SG-XPCIE1(2)FC-EM2)
- Emulex 4-Gigabit/Sec PCI Express Single and Dual FC Host Adapter (SG-XPCIE2FC-EB4-N)
- Emulex 4-Gigabit/Sec PCI Express Single and Dual FC Host Adapter (SG-XPCIE1(2)FC-EM4)
- Emulex 8-Gigabit/Sec PCI Express Single and Dual FC Host Adapter (SG-XPCIE1(2)FC-EM8-Z)
- Emulex 8-Gigabit/Sec PCI Express Single and Dual FC Host Adapter (SG-XPCIE1(2)FC-EM8-N)

**Workaround:** None.

## Guest Domain Panics When Running the cputrack Command During a Migration to a SPARC T4 System

**Bug ID 15776123:** If the cputrack command is run on a guest domain while that domain is migrated to a SPARC T4 system, the guest domain might panic on the target machine after it has been migrated.

**Workaround:** Do not run the cputrack command during the migration of a guest domain to a SPARC T4 system.

## Oracle Solaris 11: DRM Stealing Reports Oracle Solaris DR Failure and Retries

**Bug ID 15775668:** A domain that has a higher-priority policy can steal virtual CPU resources from a domain with a lower-priority policy. While this “stealing” action is in progress, you might see the following warning messages in the `ldmd` log every 10 seconds:

```
warning: Unable to unconfigure CPUs out of guest domain-name
```

**Workaround:** You can ignore these misleading messages.

## Limit the Maximum Number of Virtual Functions That Can Be Assigned to a Domain

**Bug ID 15775637:** An I/O domain has a limit on the number of interrupt resources that are available per root complex.

On SPARC T3 and SPARC T4 systems, the limit is approximately 63 MSI/X vectors. Each `igb` virtual function uses three interrupts. The `ixgbe` virtual function uses two interrupts.

If you assign a large number of virtual functions to a domain, the domain runs out of system resources to support these devices. You might see messages similar to the following:

```
WARNING: ixgbev32: interrupt pool too full.
WARNING: ddi_intr_alloc: cannot fit into interrupt pool
```

## Guest Domain That Uses Cross-CPU Migration Reports Random Uptimes After the Migration Completes

**Bug ID 15775055:** After a domain is migrated between two machines that have different CPU frequencies, the uptime reports by the `ldm list` command might be incorrect. These incorrect results occur because uptime is calculated relative to the `STICK` frequency of the machine on which the domain runs. If the `STICK` frequency differs between the source and target machines, the uptime appears to be scaled incorrectly.

The uptime reported and shown by the guest domain itself is correct. Also, any accounting that is performed by the Oracle Solaris OS in the guest domain is correct.

## Oracle Solaris 10: ixgbe Driver Might Cause a Panic When Booted With an Intel Dual Port Ethernet Controller X540 Card

**Bug ID 15773603:** When booted with an Intel dual port Ethernet Controller X540 card, the Oracle Solaris 10 `ixgbe` driver might cause a system panic. This panic occurs because the driver has a high-priority timer that blocks other drivers from attaching.

**Workaround:** Reboot the system.

## Guest Domain Console Randomly Hangs on SPARC T4 Systems

**Bug ID 15771384:** A domain's guest console might freeze if repeated attempts are made to connect to the console before and during the time the console is bound. For example, this might occur if you use an automated script to grab the console as a domain is being migrated onto the machine.

**Workaround:** To unfreeze console, perform the following commands on the domain that hosts the domain's console concentrator (usually the control domain):

```
primary# svcadm disable vntsd
primary# svcadm enable vntsd
```

## Destroying All Virtual Functions and Returning the Slots to the Root Domain Does Not Restore the Root Complex Resources

**Bug ID 15765858:** The resources on the root complex are not restored after you destroy all the virtual functions and return the slots to the root domain.

**Workaround:** Set the `io` option to `off` for the specific PCIe bus.

```
primary# ldm start-reconf primary
primary# ldm set-io iov=off pci_0
```

## ldm remove-io of PCIe Cards That Have PCIe-to-PCI Bridges Should Be Disallowed

**Bug ID 15761509:** Use only the PCIe cards that support the Direct I/O (DIO) feature, which are listed in this [support document \(https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1325454.1\)](https://support.us.oracle.com/oip/faces/secure/km/DocumentDisplay.jspx?id=1325454.1).

**Workaround:** Use the `ldm add-io` command to add the card to the primary domain again.

## ldm stop Command Might Fail If Issued Immediately After an ldm start Command

**Bug ID 15759601:** If you issue an `ldm stop` command immediately after an `ldm start` command, the `ldm stop` command might fail with the following error:

```
LDom domain stop notification failed
```

**Workaround:** Reissue the `ldm stop` command.

## init-system Does Not Restore Named Core Constraints for Guest Domains From Saved XML Files

**Bug ID 15758883:** The `ldm init-system` command fails to restore the named CPU core constraints for guest domains from a saved XML file.

**Workaround:** Perform the following steps:

1. Create an XML file for the primary domain.  

```
# ldm ls-constraints -x primary > primary.xml
```
2. Create an XML file for the guest domain or domains.  

```
# ldm ls-constraints -x ldom[,ldom][,...] > guest.xml
```
3. Power cycle the system and boot a factory default configuration.
4. Apply the XML configuration to the primary domain.  

```
# ldm init-system -r -i primary.xml
```
5. Reboot.
6. Apply the XML configuration to the guest domain or domains.  

```
# ldm init-system -f -i guest.xml
```

## System Panics When Rebooting a primary Domain That Has a Very Large Number of Virtual Functions Assigned

**Bug ID 15750727:** A system might panic when you reboot a primary domain that has a very large number of virtual functions assigned to it.

**Workaround:** Perform one of the following workarounds:

- Decrease the virtual function number to reduce the number of failed virtual functions. This change might keep the chip responsive.
- Create more Interrupt Resource Management (IRM) pools for the `ixgbe` virtual function because only one IRM pool is created by default for all the `ixgbe` virtual functions on the system.

## Partial Core primary Fails to Permit Whole-Core DR Transitions

**Bug ID 15748348:** When the primary domain shares the lowest physical core (usually 0) with another domain, attempts to set the whole-core constraint for the primary domain fail.

**Workaround:** Perform the following steps:

1. Determine the lowest bound core that is shared by the domains.  

```
# ldm list -o cpu
```
2. Unbind all the CPU threads of the lowest core from all domains other than the primary domain.  

As a result, CPU threads of the lowest core are not shared and are free for binding to the primary domain.
3. Set the whole-core constraint by doing one of the following:
  - Bind the CPU threads to the primary domain, and set the whole-core constraint by using the `ldm set -vcpu -c` command.

- Use the `ldm set -core` command to bind the CPU threads and set the whole-core constraint in a single step.

### ldm list-io Command Shows the UNK or INV State After Boot

**Bug ID 15738561:** The `ldm list -io` command might show the UNK or INV state for PCIe slots and SR-IOV virtual functions if the command runs immediately after the primary domain is booted. This problem is caused by the delay in the Logical Domains agent's reply from the Oracle Solaris OS.

This problem has been reported only on a few systems.

**Workaround:** The status of the PCIe slots and the virtual functions is automatically updated after the information is received from the Logical Domains agent.

### Migrating a Very Large Memory Domain on SPARC T4-4 Systems Results in a Panicked Domain on the Target System

**Bug ID 15731303:** Avoid migrating domains that have over 500 Gbytes of memory. Use the `ldm list -o mem` command to see the memory configuration of your domain. Some memory configurations that have multiple memory blocks that total over 500 Gbytes might panic with a stack that resembles the following:

```
panic[cpu21]/thread=2a100a5dca0:
BAD TRAP: type=30 rp=2a100a5c930 addr=6f696e740a232000 mmu_fsr=10009

sched:data access exception: MMU sfsr=10009: Data or instruction address
out of range context 0x1

pid=0, pc=0x1076e2c, sp=0x2a100a5c1d1, tstate=0x4480001607, context=0x0
g1-g7: 80000001, 0, 80a5dca0, 0, 0, 0, 2a100a5dca0

000002a100a5c650 unix:die+9c (30, 2a100a5c930, 6f696e740a232000, 10009,
2a100a5c710, 10000)
000002a100a5c730 unix:trap+75c (2a100a5c930, 0, 0, 10009, 30027b44000,
2a100a5dca0)
000002a100a5c880 unix:ktl0+64 (7022d6dba40, 0, 1, 2, 2, 18a8800)
000002a100a5c9d0 unix:page_trylock+38 (6f696e740a232020, 1, 6f69639927eda164,
7022d6dba40, 13, 1913800)
000002a100a5ca80 unix:page_trylock_cons+c (6f696e740a232020, 1, 1, 5,
7000e697c00, 6f696e740a232020)
000002a100a5cb30 unix:page_get_mnode_freelist+19c (701ee696d00, 12, 1, 0, 19, 3)
000002a100a5cc80 unix:page_get_cachelist+318 (12, 1849fe0, ffffffff, 3,
0, 1)
000002a100a5cd70 unix:page_create_va+284 (192aec0, 300ddbc6000, 0, 0,
2a100a5cf00, 300ddbc6000)
000002a100a5ce50 unix:segkmem_page_create+84 (18a8400, 2000, 1, 198e0d0, 1000,
11)
000002a100a5cf60 unix:segkmem_xalloc+b0 (30000002d98, 0, 2000, 300ddbc6000, 0,
107e290)
000002a100a5d020 unix:segkmem_alloc_vn+c0 (30000002d98, 2000, 107e000, 198e0d0,
30000000000, 18a8800)
```

```

000002a100a5d0e0 genunix:vmem_xalloc+5c8 (30000004000, 2000, 0, 0, 80000, 0)
000002a100a5d260 genunix:vmem_alloc+1d4 (30000004000, 2000, 1, 2000,
30000004020, 1)
000002a100a5d320 genunix:kmem_slab_create+44 (30000056008, 1, 300ddbc4000,
18a6840, 30000056200, 30000004000)
000002a100a5d3f0 genunix:kmem_slab_alloc+30 (30000056008, 1, ffffffff,
0, 300000560e0, 30000056148)
000002a100a5d4a0 genunix:kmem_cache_alloc+2dc (30000056008, 1, 0, b9,
ffffffffffff, 2006)
000002a100a5d550 genunix:kmem_cpucache_magazine_alloc+64 (3000245a740,
3000245a008, 7, 6028f283750, 3000245a1d8, 193a880)
000002a100a5d600 genunix:kmem_cache_free+180 (3000245a008, 6028f2901c0, 7, 7,
7, 3000245a740)
000002a100a5d6b0 ldc:vio_destroy_mblks+c0 (6028efe8988, 800, 0, 200, 19de0c0, 0)
000002a100a5d760 ldc:vio_destroy_multipools+30 (6028f1542b0, 2a100a5d8c8, 40,
0, 10, 30000282240)
000002a100a5d810 vnet:vgen_unmap_rx_dring+18 (6028f154040, 0, 6028f1a3cc0, a00,
200, 6028f1abc00)
000002a100a5d8d0 vnet:vgen_process_reset+254 (1, 6028f154048, 6028f154068,
6028f154060, 6028f154050, 6028f154058)
000002a100a5d9b0 genunix:taskq_thread+3b8 (6028ed73908, 6028ed738a0, 18a6840,
6028ed738d2, e4f746ec17d8, 6028ed738d4)

```

**Workaround:** Avoid performing migrations of domains that have over 500 Gbytes of memory.

## Removing a Large Number of CPUs From a Guest Domain Fails

**Bug ID 15726205:** You might see the following error message when you attempt to remove a large number of CPUs from a guest domain:

```

Request to remove cpu(s) sent, but no valid response received
VCPU(s) will remain allocated to the domain, but might
not be available to the guest OS
Resource modification failed

```

**Workaround:** Stop the guest domain before you remove more than 100 CPUs from the domain.

## Cannot Use Oracle Solaris Hot-Plug Operations to Hot-Remove a PCIe Endpoint Device

**Bug ID 15721872:** You cannot use Oracle Solaris hot-plug operations to hot-remove a PCIe endpoint device after that device is removed from the primary domain by using the `ldm rm-io` command. For information about replacing or removing a PCIe endpoint device, see [“Making PCIe Hardware Changes” in Oracle VM Server for SPARC 3.1 Administration Guide](#).

## nxge Panics When Migrating a Guest Domain That Has Hybrid I/O and Virtual I/O Virtual Network Devices

**Bug ID 15710957:** When a heavily loaded guest domain has a hybrid I/O configuration and you attempt to migrate it, you might see an nxge panic.

**Workaround:** Add the following line to the `/etc/system` file on the primary domain and on any service domain that is part of the hybrid I/O configuration for the domain:

```
set vsw:vsw_hio_max_cleanup_retries = 0x200
```

## All ldm Commands Hang When Migrations Have Missing Shared NFS Resources

**Bug ID 15708982:** An initiated or ongoing migration, or any `ldm` command, hangs forever. This situation occurs when the domain to be migrated uses a shared file system from another system and the file system is no longer shared.

**Workaround:** Make the shared file system accessible again.

## Logical Domains Agent Service Does Not Come Online If the System Log Service Does Not Come Online

**Bug ID 15707426:** If the system log service, `svc:/system/system-log`, fails to start and does not come online, the Logical Domains agent service will not come online. When the Logical Domains agent service is not online, the `virtinfo`, `ldm add-vsw`, `ldm add-vdsdev`, and `ldm list-io` commands might not behave as expected.

**Workaround:** Ensure that the `svc:/ldoms/agents:default` service is enabled and online:

```
# svcs -l svc:/ldoms/agents:default
```

If the `svc:/ldoms/agents:default` service is offline, verify that the service is enabled and that all dependent services are online.

## Kernel Deadlock Causes Machine to Hang During a Migration

**Bug ID 15704500:** The migration of an active guest domain might hang and cause the source machine to become unresponsive. When this problem occurs, the following message is written to the console and to the `/var/adm/messages` file:

```
vcc: i_vcc_ldc_fini: cannot close channel 15
vcc: [ID 815110 kern.notice] i_vcc_ldc_fini: cannot
close channel 15
```

Note that the channel number shown is an Oracle Solaris internal channel number that might be different for each warning message.

**Workaround:** Before you migrate the domain, disconnect from the guest domain's console.

**Recovery:** Perform a power cycle of the source machine.

## DRM and ldm list Output Shows a Different Number of Virtual CPUs Than Are Actually in the Guest Domain

**Bug ID 15702475:** A No response message might appear in the Oracle VM Server for SPARC log when a loaded domain's DRM policy expires after the CPU count has been substantially reduced. The `ldm list` output shows that more CPU resources are allocated to the domain than is shown in the `ps rinfo` output.

**Workaround:** Use the `ldm set -vcpu` command to reset the number of CPUs on the domain to the value that is shown in the `ps rinfo` output.

## Live Migration of a Domain That Depends on an Inactive Master Domain on the Target Machine Causes ldm to Fault With a Segmentation Fault

**Bug ID 15701865:** If you attempt a live migration of a domain that depends on an inactive domain on the target machine, the `ldmd` daemon faults with a segmentation fault, and the domain on the target machine restarts. Although you can still perform a migration, it will not be a live migration.

**Workaround:** Perform one of the following actions before you attempt the live migration:

- Remove the guest dependency from the domain to be migrated.
- Start the master domain on the target machine.

## DRM Fails to Restore the Default Number of Virtual CPUs for a Migrated Domain When the Policy Is Removed or Expired

**Bug ID 15701853:** After you perform a domain migration while a DRM policy is in effect, if the DRM policy expires or is removed from the migrated domain, DRM fails to restore the original number of virtual CPUs to the domain.

**Workaround:** If a domain is migrated while a DRM policy is active and the DRM policy is subsequently expired or removed, reset the number of virtual CPUs. Use the `ldm set -vcpu` command to set the number of virtual CPUs to the original value on the domain.

## Virtual CPU Timeout Failures During DR

**Bug ID 15701258:** Running the `ldm set -vcpu 1` command on a guest domain that has over 100 virtual CPUs and some cryptographic units fails to remove the virtual CPUs. The virtual CPUs are not removed because of a DR timeout failure. The cryptographic units are successfully removed.

**Workaround:** Use the `ldm rm -vcpu` command to remove all but one of the virtual CPUs from the guest domain. Do not remove more than 100 virtual CPUs at a time.

## Migration Failure Reason Not Reported When the System MAC Address Clashes With Another MAC Address

**Bug ID 15699763:** A domain cannot be migrated if it contains a duplicate MAC address. Typically, when a migration fails for this reason, the failure message shows the duplicate MAC address. However in rare circumstances, this failure message might not report the duplicate MAC address.

```
# ldm migrate ldg2 system2
Target Password:
Domain Migration of LDom ldg2 failed
```

**Workaround:** Ensure that the MAC addresses on the target machine are unique.

## Simultaneous Migration Operations in “Opposite Direction” Might Cause ldm to Hang

**Bug ID 15696986:** If two `ldm migrate` commands are issued simultaneously in the “opposite direction,” the two commands might hang and never complete. An opposite direction situation occurs when you simultaneously start a migration on machine A to machine B and a migration on machine B to machine A.

The hang occurs even if the migration processes are initiated as dry runs by using the `-n` option. When this problem occurs, all other `ldm` commands might hang.

**Workaround:** None.

## Removing a Large Number of CPUs From the Control Domain Fails

**Bug ID 15677358:** Use a delayed reconfiguration rather than dynamic reconfiguration to remove more than 100 CPUs from the control domain (also known as the primary domain). Use the following steps:

1. Use the `ldm start -reconf primary` command to put the control domain in delayed reconfiguration mode.
2. Remove the desired number of CPU resources.

If you make a mistake while removing CPU resources, do not attempt another request to remove CPUs while the control domain is still in a delayed reconfiguration. If you do so, the commands will fail (see [“Only One CPU Configuration Operation Is Permitted to Be Performed During a Delayed Reconfiguration” on page 39](#)). Instead, undo the delayed reconfiguration operation by using the `ldm cancel -reconf` command, and start over.

3. Reboot the control domain.

## System Running the Oracle Solaris 10 8/11 OS That Has the Elastic Policy Set Might Hang

**Bug IDs 15672651 and 15731467:** You might experience OS hangs at login or while executing commands when the following conditions are met:

- The Oracle Solaris 10 8/11 OS is running on a SPARC sun4v system
- The Power Management (PM) Elastic policy is set on the system's ILOM Service Processor

**Workaround:** Apply patch ID 147149-01.

## pkgadd Fails to Set ACL Entries on /var/svc/manifest/platform/sun4v/ldmd.xml

**Bug ID 15668881:** When using the pkgadd command to install the SUNWldm.v package from a directory that is exported by means of NFS from a Sun ZFS Storage Appliance, you might see the following error message:

```
cp: failed to set acl entries on /var/svc/manifest/platform/sun4v/ldmd.xml
```

**Workaround:** Ignore this message.

## SPARC T3-1: Issue With Disks That Are Accessible Through Multiple Direct I/O Paths

**Bug ID 15668368:** A SPARC T3-1 system can be installed with dual-ported disks, which can be accessed by two different direct I/O devices. In this case, assigning these two direct I/O devices to different domains can cause the disks to be used by both domains and affect each other based on the actual usage of those disks.

**Workaround:** Do not assign direct I/O devices that have access to the same set of disks to different I/O domains. To determine whether you have dual-ported disks on a SPARC T3-1 system, run the following command on the SP:

```
-> show /SYS/SASBP
```

If the output includes the following fru\_description value, the corresponding system has dual-ported disks:

```
fru_description = BD,SAS2,16DSK,LOUISE
```

If dual disks are found to be present in the system, ensure that both of the following direct I/O devices are always assigned to the same domain:

```
pci@400/pci@1/pci@0/pci@4 /SYS/MB/SASHBA0
pci@400/pci@2/pci@0/pci@4 /SYS/MB/SASHBA1
```

## Memory DR Removal Operations With Multiple Plumbed NIU nxge Instances Can Hang Indefinitely and Never Complete

**Bug ID 15667770:** When multiple NIU nxge instances are plumbed on a domain, the `ldm rm-mem` and `ldm set -mem` commands, which are used to remove memory from the domain, might never complete. To determine whether the problem has occurred during a memory removal operation, monitor the progress of the operation with the `ldm list -o status` command. You might have encountered this problem if the progress percentage remains constant for several minutes.

**Workaround:** Cancel the `ldm rm-mem` or `ldm set -mem` command, and check whether a sufficient amount of memory was removed. If not, a subsequent memory removal command to remove a smaller amount of memory might complete successfully.

If the problem has occurred on the primary domain, do the following:

1. Start a delayed reconfiguration operation on the primary domain.  

```
# ldm start-reconf primary
```
2. Assign the desired amount of memory to the domain.
3. Reboot the primary domain.

If the problem occurred on another domain, stop the domain before adjusting the amount of memory that is assigned to the domain.

## Using the `ldm stop -a` Command on Domains in a Master-Slave Relationship Leaves the Slave With the Stopping Flag Set

**Bug ID 15664666:** When a reset dependency is created, an `ldm stop -a` command might result in a domain with a reset dependency being restarted instead of only stopped.

**Workaround:** First, issue the `ldm stop` command to the master domain. Then, issue the `ldm stop` command to the slave domain. If the initial stop of the slave domain results in a failure, issue the `ldm stop -f` command to the slave domain.

## Migration of a Domain That Has an Enabled Default DRM Policy Results in a Target Domain Being Assigned All Available CPUs

**Bug ID 1565513:** Following the migration of an active domain, CPU utilization in the migrated domain can increase dramatically for a short period of time. If a dynamic resource management (DRM) policy is in effect for the domain at the time of the migration, the Logical Domains Manager might begin to add CPUs. In particular, if the `vcpu-max` and `at tack` properties were not specified when the policy was added, the default value of `unlimited` causes all the unbound CPUs in the target machine to be added to the migrated domain.

**Recovery:** No recovery is necessary. After the CPU utilization drops below the upper limit that is specified by the DRM policy, the Logical Domains Manager automatically removes the CPUs.

## An In-Use MAC Address Can be Reassigned

**Bug ID 15655199:** Sometimes an in-use MAC address is not detected and it is erroneously reassigned.

**Workaround:** Manually ensure that an in-use MAC address cannot be reassigned.

## ldmconfig Cannot Create a Domain Configuration on the SP

**Bug ID 15654965:** The `ldmconfig` script cannot properly create a stored domain configuration on the service processor (SP).

**Workaround:** Do not power cycle the system after the `ldmconfig` script completes and the domain reboots. Instead, perform the following manual steps:

1. Add the configuration to the SP.
 

```
# ldm add-spconfig new-config-name
```
2. Remove the `primary-with-clients` configuration from the SP.
 

```
# ldm rm-spconfig primary-with-clients
```
3. Power cycle the system.

If you do not perform these steps prior to the system being power cycled, the existence of the `primary-with-client` configuration causes the domains to be inactive. In this case, you must bind each domain manually and then start them by running the `ldm start -a` command. After the guests have booted, repeating this sequence enables the guest domains to be booted automatically after a power cycle.

## Uncooperative Oracle Solaris Domain Migration Can Be Blocked If `cpu0` Is Offline

**Bug ID 15653424:** The migration of an active domain can fail if it is running a release earlier than the Oracle Solaris 10 10/09 OS release and the lowest-numbered CPU in the domain is in the `offline` state. The operation fails when the Logical Domains Manager uses CPU DR to reduce the domain to a single CPU. In doing so, the Logical Domains Manager attempts to remove all but the lowest CPU in the domain but because that CPU is offline, the operation fails.

**Workaround:** Before attempting the migration, ensure that the lowest-numbered CPU in the domain is in the `online` state.

## Memory DR Is Disabled Following a Canceled Migration

**Bug ID 15646293:** After an Oracle Solaris 10 9/10 domain has been suspended as part of a migration operation, memory dynamic reconfiguration (DR) is disabled. This action occurs not only when the migration is successful but also when the migration has been canceled, despite the fact that the domain remains on the source machine.

## Dynamic Reconfiguration of MTU Values of Virtual Network Devices Sometimes Fails

**Bug ID 15631119:** If you modify the maximum transmission unit (MTU) of a virtual network device on the control domain, a delayed reconfiguration operation is triggered. If you subsequently cancel the delayed reconfiguration, the MTU value for the device is not restored to the original value.

**Recovery:** Rerun the `ldm set -vnet` command to set the MTU to the original value. Resetting the MTU value puts the control domain into delayed reconfiguration mode, which you need to cancel. The resulting MTU value is now the original, correct MTU value.

```
# ldm set-vnet mtu=orig-value vnet1 primary
# ldm cancel-op reconf primary
```

## Migrated Domain With MAUs Contains Only One CPU When Target OS Does Not Support DR of Cryptographic Units

**Bug ID 15606220:** Starting with the Logical Domains 1.3 release, a domain can be migrated even if it has one or more cryptographic units bound to it.

In the following circumstances, the target machine will contain only one CPU after the migration is completed:

- Target machine runs Logical Domains 1.2
- Control domain on the target machine runs a version of the Oracle Solaris OS that does not support cryptographic unit DR
- You migrate a domain that contains cryptographic units

After the migration completes, the target domain will resume successfully and be operational, but will be in a degraded state (just one CPU).

**Workaround:** Prior to the migration, remove the cryptographic unit or units from the source machine that runs Logical Domains 1.3.

**Mitigation:** To avoid this issue, perform one or both of these steps:

- Install the latest Oracle VM Server for SPARC software on the target machine.
- Install patch ID 142245-01 on the control domain of the target machine, or upgrade to at least the Oracle Solaris 10 10/09 OS.

## Confusing Migration Failure Message for Real Address Memory Bind Failures

**Bug ID 15605806:** In certain situations, a migration fails with the following error message, and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This

situation can occur even if the total amount of available memory on the target machine is greater than the amount of memory being used by the source domain (as shown by `ldm ls-devices -a mem`).

```
Unable to bind 29952M memory region at real address 0x8000000
Domain Migration of LDom ldg0 failed
```

**Cause:** This failure is due the inability to meet congruence requirements between the Real Address (RA) and the Physical Address (PA) on the target machine.

**Workaround:** Stop the domain and perform the migration as a cold migration. You can also reduce the size of the memory on the guest domain by 128 Mbytes, which might permit the migration to proceed while the domain is running.

## Dynamically Removing All the Cryptographic Units From a Domain Causes SSH to Terminate

**Bug ID 15600969:** If all the hardware cryptographic units are dynamically removed from a running domain, the cryptographic framework fails to seamlessly switch to the software cryptographic providers, and kills all the ssh connections.

**Recovery:** Re-establish the ssh connections after all the cryptographic units are removed from the domain.

**Workaround:** Set `UseOpenSSLEngine=no` in the `/etc/ssh/sshd_config` file on the server side, and run the `svcadm restart ssh` command.

All ssh connections will no longer use the hardware cryptographic units (and thus not benefit from the associated performance improvements), and ssh connections will not be disconnected when the cryptographic units are removed.

## PCI Express Dual 10-Gigabit Ethernet Fiber Card Shows Four Subdevices in `ldm list-io -l` Output

**Bug ID 15597025:** When you run the `ldm ls-io -l` command on a system that has a PCI Express Dual 10-Gigabit Ethernet Fiber card (X1027A-Z) installed, the output might show the following:

```
primary# ldm ls-io -l
...
pci@500/pci@0/pci@c PCIE5 OCC primary
network@0
network@0,1
ethernet
ethernet
```

The output shows four subdevices even though this Ethernet card has only two ports. This anomaly occurs because this card has four PCI functions. Two of these functions are disabled internally and appear as ethernet in the `ldm ls-io -l` output.

**Workaround:** You can ignore the ethernet entries in the `ldm ls -io -l` output.

## Using Logical Domains mpgroup With MPXIO Storage Array Configuration for High-Disk Availability

**Bug ID 15591769:** When creating a LUN, you can add it to the virtual disk service for both primary and alternate domains by using the same mpgroup. To specify which domain to use first when accessing the LUN, add that virtual disk service device first.

- To use the LUN from `primary-vds0` first, perform the following commands:

```
primary# ldm add-vdsdev mpgroup=ha lun1@primary-vds0
primary# ldm add-vdsdev mpgroup=ha lun1@alternate-vds0
primary# ldm add-vdisk disk1 lun1@primary-vds0 gd0
```

- To use the LUN from `alternate-vds0` first, perform the following commands:

```
primary# ldm add-vdsdev mpgroup=ha lun1@alternate-vds0
primary# ldm add-vdsdev mpgroup=ha lun1@primary-vds0
primary# ldm add-vdisk disk1 lun1@alternate-vds0 gd0
```

## ldm Commands Are Slow to Respond When Several Domains Are Booting

**Bug ID 15572184:** An `ldm` command might be slow to respond when several domains are booting. If you issue an `ldm` command at this stage, the command might appear to hang. Note that the `ldm` command will return after performing the expected task. After the command returns, the system should respond normally to `ldm` commands.

**Workaround:** Avoid booting many domains simultaneously. However, if you must boot several domains at once, refrain from issuing further `ldm` commands until the system returns to normal. For instance, wait for about two minutes on Sun SPARC Enterprise T5140 and T5240 servers and for about four minutes on the Sun SPARC Enterprise T5440 server or Sun Netra T5440 server.

## Oracle Solaris 11: Zones Configured With an Automatic Network Interface Might Fail to Start

**Bug ID 15560811:** In Oracle Solaris 11, zones that are configured with an automatic network interface (`anet`) might fail to start in a domain that has Logical Domains virtual network devices only.

- **Workaround 1:** Assign one or more physical network devices to the guest domain. Use PCIe bus assignment, the Direct I/O (DIO), or the SR-IOV feature to assign a physical NIC to the domain.
- **Workaround 2:** If the zones configuration requirement is to have interzone communication only within the domain, create an `etherstub` device. Use the `etherstub` device as the “lower link” in the zones configuration so that that virtual NICs are created on the `etherstub` device.

- Workaround 3:** Use exclusive link assignment to assign a Logical Domains virtual network device to a zone. Assign virtual network devices, as needed, to the domain. You might also choose to disable inter-vnet links to be able to create a large number of virtual network devices.

## Oracle Solaris 10: Virtual Network Devices Are Not Created Properly on the Control Domain

**Bug ID 15560201:** Sometimes `ifconfig` indicates that the device does not exist after you add a virtual network or virtual disk device to a domain. This situation might occur as the result of the `/devices` entry not being created.

Although this problem should not occur during normal operation, the error sometimes occurs when the instance number of a virtual network device does not match the instance number listed in `/etc/path_to_inst` file.

For example:

```
# ifconfig vnet0 plumb
ifconfig: plumb: vnet0: no such interface
```

The instance number of a virtual device is shown under the `DEVICE` column in the `ldm list` output:

```
# ldm list -o network primary
NAME
primary

MAC
00:14:4f:86:6a:64

VSW
NAME      MAC          NET-DEV  DEVICE  DEFAULT-VLAN-ID  PVID  VID  MTU  MODE
primary-vsw0 00:14:4f:f9:86:f3 nxge0    switch@0 1              1      1500

NETWORK
NAME  SERVICE          DEVICE  MAC          MODE  PVID  VID  MTU
vnet1 primary-vsw0@primary network@0 00:14:4f:f8:76:6d 1      1500
```

The instance number (0 for both the `vnet` and `vsw` shown previously) can be compared with the instance number in the `path_to_inst` file to ensure that they match.

```
# egrep '(vnet|vsw)' /etc/path_to_inst
"/virtual-devices@100/channel-devices@200/virtual-network-switch@0" 0 "vsw"
"/virtual-devices@100/channel-devices@200/network@0" 0 "vnet"
```

**Workaround:** In the case of mismatching instance numbers, remove the virtual network or virtual switch device. Then, add them again by explicitly specifying the instance number required by setting the `id` property.

You can also manually edit the `/etc/path_to_inst` file. See the [path\\_to\\_inst\(4\)](#) man page.




---

**Caution** – Changes should not be made to `/etc/path_to_inst` without careful consideration.

---

## Newly Added NIU/XAUI Adapters Are Not Visible to the Host OS If Logical Domains Is Configured

**Bug ID 1555509:** When Logical Domains is configured on a system and you add another XAUI network card, the card is not visible after the machine has undergone a power cycle.

**Recovery:** To make the newly added XAUI visible in the control domain, perform the following steps:

1. Set and clear a dummy variable in the control domain.

The following commands use a dummy variable called `fix-xaui`:

```
# ldm set-var fix-xaui=yes primary
# ldm rm-var fix-xaui primary
```

2. Save the modified configuration to the SP, replacing the current configuration.

The following commands use a configuration name of `config1`:

```
# ldm rm-spconfig config1
# ldm add-spconfig config1
```

3. Perform a reconfiguration reboot of the control domain.

```
# reboot -- -r
```

At this time, you can configure the newly available network or networks for use by Logical Domains.

## I/O Domain or Guest Domain Panics When Booting From e1000g

**Bug ID 15543982:** You can configure a maximum of two domains with dedicated PCI-E root complexes on systems such as the Sun Fire T5240. These systems have two UltraSPARC T2 Plus CPUs and two I/O root complexes.

`pci@500` and `pci@400` are the two root complexes in the system. The primary domain will always contain at least one root complex. A second domain can be configured with an unassigned or unbound root complex.

The `pci@400` fabric (or leaf) contains the on-board `e1000g` network card. The following circumstances could lead to a domain panic:

- If the system is configured with a primary domain that contains `pci@500` and a second domain that contains `pci@400`

---

**Note** – For some blades, the primary domain (system disk) is on the pci@400 bus by default.

---

- The e1000g device on the pci@400 fabric is used to boot the second domain

Avoid the following network devices if they are configured in a non-primary domain:

```
/pci@400/pci@0/pci@c/network@0,1
/pci@400/pci@0/pci@c/network@0
```

When these conditions are true, the domain will panic with a PCI-E Fatal error.

Avoid such a configuration or, if the configuration is used, do not boot from the listed devices.

## Explicit Console Group and Port Bindings Are Not Migrated

**Bug ID 15527921:** During a migration, any explicitly assigned console group and port are ignored, and a console with default properties is created for the target domain. This console is created using the target domain name as the console group and using any available port on the first virtual console concentrator (vcc) device in the control domain. If there is a conflict with the default group name, the migration fails.

**Recovery:** To restore the explicit console properties following a migration, unbind the target domain and manually set the desired properties using the `ldm set -vcons` command.

## Migration Does Not Fail If a vdsdev on the Target Has a Different Back End

**Bug ID 15523133:** If the virtual disk on the target machine does not point to the same disk back end that is used on the source machine, the migrated domain cannot access the virtual disk using that disk back end. A hang can result when accessing the virtual disk on the domain.

Currently, the Logical Domains Manager checks only that the virtual disk volume names match on the source and target machines. In this scenario, no error message is displayed if the disk back ends do not match.

**Workaround:** When you are configuring the target domain to receive a migrated domain, ensure that the disk volume (vdsdev) matches the disk back end used on the source domain.

**Recovery:** Do one of the following if you discover that the virtual disk device on the target machine points to the incorrect disk back end:

- Migrate the domain and fix the vdsdev.
  1. Migrate the domain back to the source machine.
  2. Fix the vdsdev on the target to point to the correct disk back end.
  3. Migrate the domain to the target machine again.

- Stop and unbind the domain on the target, and fix the vdsdev. If the OS supports virtual I/O dynamic reconfiguration and the incorrect virtual disk is not in use on the domain (that is, it is not the boot disk and is unmounted), do the following:
  1. Use the `ldm rm-vdisk` command to remove the disk.
  2. Fix the vdsdev.
  3. Use the `ldm add-vdisk` command to add the virtual disk again.

## Migration Can Fail to Bind Memory Even If the Target Has Enough Available

**Bug ID 15523120:** In certain situations, a migration fails and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This situation can occur even if the total amount of available memory on the target machine is greater than the amount of memory being used by the source domain.

This failure occurs because migrating the specific memory ranges in use by the source domain requires that compatible memory ranges are available on the target as well. When no such compatible memory range is found for any memory range in the source, the migration cannot proceed.

**Recovery:** If this condition is encountered, you might be able to migrate the domain if you modify the memory usage on the target machine. To do this, unbind any bound or active logical domain on the target.

Use the `ldm list-devices -a mem` command to see what memory is available and how it is used. You might also need to reduce the amount of memory that is assigned to another domain.

## Logical Domains Manager Does Not Start If the Machine Is Not Networked and an NIS Client Is Running

**Bug ID 15518409:** If you do not have a network configured on your machine and have a Network Information Services (NIS) client running, the Logical Domains Manager will not start on your system.

**Workaround:** Disable the NIS client on your non-networked machine:

```
# svcadm disable nis/client
```

## Logical Domains Manager Displays Migrated Domains in Transition States When They Are Already Booted

**Bug ID 15516245:** On occasion, an active logical domain appears to be in the *transition* state instead of the *normal* state long after it is booted, or following the completion of a domain migration. This glitch is harmless, and the domain is fully operational. To see what flag is set, check the `flags` field in the `ldm list -l -p` command output, or check the `FLAGS` field in the `ldm list` command, which shows `-n----` for *normal* or `-t----` for *transition*.

**Recovery:** After the next reboot, the domain shows the correct state.

## Cannot Connect to Migrated Domain's Console Unless vntsd Is Restarted

**Bug ID 15513998:** Occasionally, after a domain has been migrated, it is not possible to connect to the console for that domain.

**Workaround:** Restart the vntsd SMF service to enable connections to the console:

```
# svcadm restart vntsd
```

---

**Note** – This command will disconnect all active console connections.

---

## Sometimes, Executing the uadmin 1 0 Command From a Logical Domains System Does Not Return the System to the OK Prompt

**Bug ID 15511551:** Sometimes, executing the `uadmin 1 0` command from the command line of a Logical Domains system does not leave the system at the ok prompt after the subsequent reset. This incorrect behavior is seen only when the Logical Domains variable `auto-reboot?` is set to `true`. If `auto-reboot?` is set to `false`, the expected behavior occurs.

**Workaround:** Use this command instead:

```
uadmin 2 0
```

Or, always run with `auto-reboot?` set to `false`.

## Logical Domains Manager Can Take Over 15 Minutes to Shut Down a Domain

**Bug ID 15505014:** A domain shutdown or memory scrub can take over 15 minutes with a single CPU and a very large memory configuration. During a shutdown, the CPUs in a domain are used to scrub all the memory owned by the domain. The time taken to complete the scrub can be quite long if a configuration is imbalanced, for example, a single CPU domain with 512 Gbytes of memory. This prolonged scrub time extends the amount of time needed to shut down a domain.

**Workaround:** Ensure that large memory configurations (more than 100 Gbytes) have at least one core.

## scadm Command Can Hang Following an SC or SP Reset

**Bug ID 15469227:** The `scadm` command on a control domain running at least the Oracle Solaris 10 5/08 OS can hang following an SC reset. The system is unable to properly re-establish a connection following an SC reset.

**Recovery:** Reboot the host to re-establish connection with the SC.

## Simultaneous Net Installation of Multiple Domains Fails When in a Common Console Group

**Bug ID 15453968:** Simultaneous net installation of multiple guest domains fails on systems that have a common console group.

**Workaround:** Only net-install on guest domains that each have their own console group. This failure is seen only on domains with a common console group shared among multiple net-installing domains.

## Guest Domain With Too Many Virtual Networks on the Same Network Using DHCP Can Become Unresponsive

**Bug ID 15422900:** If you configure more than four virtual networks (vnets) in a guest domain on the same network using the Dynamic Host Configuration Protocol (DHCP), the guest domain can eventually become unresponsive while running network traffic.

**Workaround:** Set `ip_ire_min_bucket_cnt` and `ip_ire_max_bucket_cnt` to larger values, such as 32, if you have 8 interfaces.

**Recovery:** Issue an `ldm stop-domain ldom` command followed by an `ldm start-domain ldom` command on the guest domain (*ldom*) in question.

## OpenBoot PROM Variables Cannot be Modified by the eeprom Command When the Logical Domains Manager Is Running

**Bug ID 15387338:** This issue is summarized in [“Logical Domains Variable Persistence” on page 34](#) and affects only the control domain.

## Cannot Set Security Keys With Logical Domains Running

**Bug ID 15370442:** The Logical Domains environment does not support setting or deleting wide-area network (WAN) boot keys from within the Oracle Solaris OS by using the `ickey(1M)` command. All `ickey` operations fail with the following error:

```
ickey: setkey: ioctl: I/O error
```

In addition, WAN boot keys that are set using OpenBoot firmware in logical domains other than the control domain are not remembered across reboots of the domain. In these domains, the keys set from the OpenBoot firmware are valid only for a single use.

## Behavior of the ldm stop-domain Command Can Be Confusing

**Bug ID 15368170:** In some cases, the behavior of the `ldm stop-domain` command is confusing.

```
# ldm stop-domain -f ldom
```

If the domain is at the kernel module debugger, `kldb(1)`, prompt, then the `ldm stop-domain` command fails with the following error message:

```
LDom <domain-name> stop notification failed
```

## Documentation Issues

This section contains documentation issues and errors that have been found too late to resolve for the Oracle VM Server for SPARC 3.1 release.

### **ldm(1M) Man Page: Describe the Limitation for Using the mblock Property**

**Bug ID 18105821:** The `ldm(1M)` man page does not clearly describe the limitation that you cannot use the `mblock` property to specify the physical addresses of DIMMs.

### **ldm(1M) Man Page: Improve description of the ldm list -o status Command**

**Bug ID 17796758:** The `ldm(1M)` man page does not clearly describe the `ldm list -o status` output. You can use the `-o status` option to show the status of any migration operations or DR operations that are in progress. This information is derived from the flags in the `FLAGS` field. The `-o status` option does not relate to the `STATE` field.

### **ldm(1M) Man Page: Only ldm add-spconfig -r Performs a Manual Recovery**

The description of the `-r` option in the `ldm(1M)` man page currently states that the `add-spconfig`, `list-spconfig`, and `remove-spconfig` subcommands use this option to perform a manual recovery. This is not true. Only the `ldm add-spconfig -r` command can be used to perform the manual recovery.

### **Oracle VM Server for SPARC 3.1 Administration Guide Fibre Channel SR-IOV OS Requirements Are Incorrect**

To use the SR-IOV feature, all domains must run at least the Oracle Solaris 11.1.17.4.0 OS or the Oracle Solaris 10 1/13 OS plus Oracle VM Server for SPARC 3.1.1 patch ID 150817-02 and the required patches in [Table 1–2](#).

## Resolved Issues

### Resolved Issues in the Oracle VM Server for SPARC 3.1.1.2 Release

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.1.1.2 software release:

- 19864344    Fix problem with SSLv3 for CVE-2014-3566
- 17889357    ldm set-vnet not updating vnet\_node alt-mac-addr list
- 17797079    vnet\_del\_alt\_mac\_addrs() prints a warning msg to SMF
- 17761714    ldm add-vcpu fails with Power Management Failure
- 17387444    ldomVnetChange trap is not sent when alt-mac-addr is changed

### Resolved Issues in the Oracle VM Server for SPARC 3.1.1.1 Release

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.1.1.1 software release:

- 19480835    Very large LDC counts can result in Solaris issues in guests
- 19238315    Found memory leak in pvlan config validation function.
- 18895023    Problem in MGMTLDMGR/LDOMMGR-MGMT
- 18726175    ldmd aborts when adding memory to guest domain
- 18595023    SR-IOV cannot be used on PCI-BOX SLOT#4 or higher
- 18594819    Enable dynamic PCIe slots on Fujitsu M10
- 18479243    ldmp2v\_prepare failing on large partition
- 18477335    HV Abort on deleteboard unbind=resource
- 17934416    Cold migration "Failed to read feasibility response type (9)"
- 17796639    ldmd core dump when attempting to bind domain in BINDING/UNBINDING state

## Resolved Issues in the Oracle VM Server for SPARC

### 3.1.1 Release

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.1.1 software release:

- 15584929 Add Logical Domains Manager support to config network bandwidth limit
- 15726854 RFE: ldmd executes numerous misaligned loads during a migration
- 15738764 RFE: VM-API support for Oracle VM Server for SPARC
- 15748028 RFE: ldmd should send an XML event when an OBP variable or VM-API property is modified
- 15814176 RFE: Need a way to create and destroy multiple VFs in one command
- 16494899 Received invalid physical address during migration of domain
- 16922112 Logical Domains Manager core dump if disabled ldmd when there is a domain in a transition state
- 17179630 Oracle VM Server for SPARC support for Fibre Channel SR-IOV
- 17188920 ldmpower's options suppress and timestamps are displaying nothing
- 17372658 Memory leak in vnet\_list\_one()
- 17478218 Fibre Channel create-vf fails when primary's main network port is down
- 17596585 ldmpower does not accommodate guest domain names longer than 13 characters
- 17694771 After migration virtual network won't init: exceeded number of permitted handshake attempts
- 17707801 Hypervisor needs to be aware of guest domain cpu-arch=sparc64-class1
- 17732946 Disable memory blacklisting
- 17742095 sparc64-class1 migration failed since generic CPU version not set
- 17768672 Attempt to destroy multiple virtual functions crashes if referenced physical function has zero virtual functions
- 17777004 Remove degraded memory prior to Board DR/live migration
- 17777212 Board DR core dump on attempt to remove unassigned memory region
- 17960227 ldmd core dump on Oracle Solaris 10 during Board DR
- 18055802 English and Japanese man pages should be updated for 3.1.1
- 18112775 OVM core dump on set-mem/rm-mem of Oracle Solaris 10 domain

- 18112822 ldm core dump on Oracle Solaris 10 during Board DR
- 18115873 Migration of guest running pre Oracle Solaris 10 1/13 OS from 3.0 to 3.1 will generate STICK err

## Resolved Issues in the Oracle VM Server for SPARC

### 3.1.0.1 Release

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.1.0.1 software release:

- 16958880 ldm stop-domain -t doesn't work as the man page explains
- 17043095 null-pointer-deref in parse\_mac\_addr\_list of iov\_
- 17043143 use-after-free in client\_loop of clients.c:735
- 17215630 Support for Fujitsu Board DR
- 17245915 Resizing guest's vcpu after core fault coredump
- 17285385 Need to enable NPRD support for Fujitsu M10
- 17335156 Integrate FJ Board DR Support
- 17511365 Promote FJ Board DR Development Feature To Product Feature
- 17621771 OS panic after Board DR remaps degraded memory
- 17635306 HV local memory not assigned

## Resolved Issues in the Oracle VM Server for SPARC 3.1

### Release

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.1 software release:

- 7151727 cancel-reconf does not remove DRM policy plus other values
- 7151870 Japanese man pages need to be added to the 2.2 IPS manifests
- 7152717 LDoms man pages all refer to "SunOS 5.10" (even on S11)
- 7160611 Package system/ldoms/ldomsmanager delivers broken symlink
- 7192748 Stop CMU fails when control domain suspend times out prematurely
- 7198319 ldm abort when restoring previously evacuated root complexes

- 
- 7200165 ldmpower shows 'MIB communication error' on unsupported platforms
  - 7201404 Restore ldmpower to SVR4 and IPS packaging
  - 7202777 Add Oracle VM Server for SPARC 3.0 man pages
  - 7202837 Crash in core\_remove() if vcpu unconfigure fails
  - 7203080 Core remap is not attempted during DR evacuation in some cases
  - 7203749 Logical Domains Manager core dump during migration when PM policy is switched from perf to elastic
  - 7203865 Logical Domains Manager should set next poweron config after recovery
  - 7204068 Need to cleanup unused vars in pm\_ops
  - 7204705 virt\_ops should be statically set to ldm\_ops
  - 7204732 ldmd dumped core on a set-mem command
  - 7205057 IOV: Inactive guest allows add-io of root complex + virtual function from diff bus + bind and start
  - 7205900 Guest domain fails to bootup when the memory is highly fragmented
  - 7206202 Truncated hostid printed by ldm list commands
  - 7206216 DIO device is gone after reboot
  - 7206310 HV dump collection fails due to use of 32-bit statvfs
  - 14851589 Need to clean up Fujitsu M10 system descriptions in Japanese and English man pages
  - 14851668 Oracle VM Server for SPARC must clear HV dump memory after dump collection
  - 14851717 Memory should be in pre-charged powerdown mode when elastic mode is set
  - 14851768 cancel-reconf triggers Assertion failed: pcpup->res\_id != 0
  - 14851788 cancel-reconf triggers fatal error: HV MD hv\_mblock ... has invalid resource\_id
  - 14851800 Hypervisor dump does not save HV memory base addresses
  - 14851813 ldmd crashes when destroying and rebuilding domain
  - 15509749 Migration hv code should go through the backend ops structure
  - 15517293 XML parser errors are logged in ldmd SVC log during migration
  - 15605806 Migration failure message should be improved for real address memory
  - 15698622 Assertion failure in cleanup\_suspend\_failure()

- 15712764 PM code cleanup: better variable names, code organization
- 15714059 RFE: inactive domains should display the system MAC address
- 15720542 pm\_pap\_sat-1.0\_lib.so causes wsdiff to report a false positive
- 15738370 ldmd fails to reconnect DS port after cancelling Live Migration
- 15748555 SRIOV: Error message can be improved to indicate reason for failure
- 15770772 Cold migration ignores ra\_seg\_alloc() assumption it is on the seq
- 15780356 RFE: Alternate mac-address allocation support for vnet devices
- 15781142 ldm create silently accepts any unsupported command-line options
- 15785203 Unnecessary context bits restriction
- 15785454 mdstore DS timeout handling should be improved
- 15789756 Remove all references to PM\_MANAGED, as this state no longer exists
- 15791937 Checks preventing elastic policy are not needed
- 15793876 Remove benchmark from LDOMs papsat library
- 15794575 DR\_VIO\_MAX\_ERR\_SIZE should be increased
- 15798826 In ldmd XML response primary domain's vcons property is inconsistent
- 15801579 HVctl\_op\_guest\_stop fails if domain reboots in a tight loop
- 15805135 RFE: Cross-CPU migration support for SPARC M5/SPARC T5 based platforms
- 15805441 PM policy constraints should be platform specific
- 15807703 Various minor memory leaks occur during migration
- 15809205 Enhance msg on x-cpu migration error from 8.4 firmware to 8.3 or before
- 15810663 lint of IPS pkgs fails after clobber because .mog files don't exist
- 15810915 The PM policy applied to LDOMs managed resources should be at the guest granularity when possible
- 15810998 pwracap.py fails stating: Not all processor PM states were lowered
- 15811098 Evacuated I/O devices need to be shown via the CLI
- 15813210 RFE: Enable ASLR (Address Space Layout Randomization) on LDOMs code
- 15813782 Warning: MAC address for virtual function not found in virtual function mac list
- 15814289 Fujitsu M10 system Core needs to support Board DR
- 15814619 Resume timeout for a suspended guest might need to be increased

- 15814620 Logical Domains Manager crash from rolling back a memory remap that HV completed
- 15814714 Duplicate resources should be filtered out of board DR requests
- 15815064 Need a correct messages when LDoms assigns the USB storage devices
- 15815183 All incoming resources should be added to a single domain PDom
- 15815958 HV dump information must be updated during Board DR
- 15815995 Abort seen when guest domain suspend fails during CMU DR operation
- 15817307 Observability mod comments needs cleaning up
- 15817449 ldmd abort during CMU remove when the host has bound domains
- 15818190 ldm ls-io output is out of order after CMU rm then add
- 15818302 RFE: ldm rm-io should accept --dry-run as an alias for -n
- 15818483 NULL pointer reference in var\_config\_backup\_forward()
- 15818734 Setting vid argument for a virtual function can crash Logical Domains Manager for physical function device
- 15819126 Recovery Mode should recreate Direct I/O and SR-IOV configs
- 15819829 Incorrect description of papsat residency output in log file
- 15820119 DR sequencer goes into an infinite loop when vrc removal fails
- 15822221 Recovery mode should be fully automatic and not require a powercycle
- 15823584 Leaks in PM code
- 15824064 DIOV failures during board DR need more helpful messages
- 15825060 Panic in dev\_mondo queue configuration after cross CPU migration
- 15825259 Board DR: Support blacklist without evacuation
- 15825330 LDOMS\_3.0 SPARC T4-4: unable to bind guest domain
- 15825900 Adding virtual function to guest failed
- 15825992 CMU remove assert: !((ldp)->mem\_dr\_state.flag.add == true)
- 15827166 ldmd core dumped cancel-reconf
- 15829698 Fix for 7169681 precludes unmodified factory-default config: breaks ldmconfig
- 15846754 PM policy var update causes [newer] autosave for booted config
- 15847279 Logical Domains Manager did not remove the req mem when mem is available to be removed

- 15850133 PM should base PM state update decisions on req\_state field
- 15852479 Error message for system controller out of memory need correction
- 15853660 PM needs to always send the calculated new state for a resource to HV
- 15858713 ldmd aborts if all memory in a domain is degraded
- 15858717 CoD free CPU resource calculation is incorrect when a strand is degraded
- 15858722 CoD framework skips affinity routine if there are no extra free permits
- 15858731 Uncooperative migrations with Fujitsu M10 system should not be allowed
- 15861940 LDoms spconfig reports current booted config as Degraded when its not degraded
- 15872566 ldm ls -l displays empty HOSTID section for inactive guests
- 15873266 ldmd in maintenance mode at boot due to “Failed to send command to HV”
- 15873633 ldmd dumps core with a SEGV in PM Observability
- 15887526 Cancel delayed reconfig causes ldmd crash with HV MD hv\_mblock invalid message
- 15892643 XML list-bindings returns stale CPU utilization
- 15899929 Restarting Logical Domains Manager can cause loss of IOV capability in RC constraints
- 15929827 CHAIN\_ITER needs replacing with CHAIN\_ITER\_VOL in pm\_policy\_via\_ds.c
- 15942029 Permanent memory evacuation causes ldmd abort when trying to get response text
- 15962480 The guest domain state was left “stopping” after the domain was rebooted
- 15963127 Timestamps (-t) are not getting displayed with component (-c) option
- 15968595 Blacklist evacuation does not handle invalid CPU resource IDs
- 15973628 Merge Fujitsu PM code into base Oracle PM code
- 15979466 epm\_create\_node is too unwieldy to debug with via mdb
- 15981073 All ldm ls variant CLIs do not display alt-mac-addr in the -p output
- 15982523 Warning: Memory allocation of 64G is below the minimum
- 15983947 LDoms Derived PM Policy wrongly set to ELASTIC (elastic + perf = perf)
- 15986907 Alt-mac-addr are always explicitly represented in XML files with values
- 15991764 iov\_vrc\_remove() uses bio variable after it is freed

- 
- 15997828 Modify LDoms-PM code to handle new compact PRI format
  - 15998505 DR of a board with multiple domains fails with ETIME in hv\_susrec\_async\_check
  - 16004946 With multiple vnets using alt-mac-addr, XML recreation always returns failure
  - 16005218 ldm bind failed due to ASSERT(zeus.reload\_active\_state); hit in frag\_free()
  - 16006042 Physical function/virtual function status:INV, after non-primary root domain was booted up
  - 16006055 Hide InfiniBand physical function if the system FW does not support DIOV (ver < 8.4.0)
  - 16006078 Removing PCIe endpoint device hosting virtual function assigned to guest causes ldmd coredump
  - 16011596 ldmpower's message could be clear when the user does not have authority
  - 16012812 Assertion failed: state->progress <= 100, file dr\_mem.c, line 341
  - 16016576 Blacklisted resources reported evacuation pending, subsequent to reboot
  - 16017526 Guest domain with S10 OS not in lowest cycle skip state when idle
  - 16019053 set-vcc requires port-range=2000-2015 to bind range port-range=2000-2014
  - 16036742 Logical Domains Manager hits ASSERT on stop /SYS/CMU3 on SPARC T5
  - 16038554 Failed to evac and blacklist faulty memory row from guest domain
  - 16042429 "Unable to parse bootconfig HV MD" seen on add-spconfig -r config boot
  - 16050372 om\_new\_dom\_list should be protected by a mutex
  - 16052643 ldm add-io to an ldom at OBP should display a valid error message
  - 16062148 Parfait 1.1 finds some memory leaks and a NULL ptr dereference
  - 16062179 Misc lint cleanup and eradicate E\_SUSPICIOUS\_COMPARISON & E\_FUNC\_ARG\_UNUSED
  - 16070919 LDoms error message is inappropriate while performing SRIOV
  - 16074693 LDoms PM is querying the SP via SNMP too aggressively
  - 16087954 Guest domain was bound without a console
  - 16094241 Shared processors in S10 domains choose same power level for all cores
  - 16105461 Board DR does not implement full progress heartbeat protocol
  - 16164394 Getting resource PM state should support single state resources
  - 16165038 LDoms-PM Obs Mod is not closing sockets; ldmd runs out of file descriptors

- 16166749 ldm core iov\_gen\_pfvf\_drv\_props() when ldm set-io unicast-slots=<above\_allowed>
- 16166910 ldm init-system -r -I file doesn't properly rebuild split bus config
- 16171044 MAC string parsing needs to be consolidated, optimized and free of strcpy() calls
- 16172239 Core evacuation completed, core status does not get blacklisted
- 16172976 Logical Domains Manager crashes on list-io -l command when target is InfiniBand physical function device
- 16173609 ldm ls-io is out of order after DCU add
- 16178876 Live migration on SPARC T5-2 drops pm-rm messages sent from domain being migrated
- 16187066 LDoms should only apply new power cap adjustments
- 16198869 System seems to be sending a sigabrt explicitly, from hv calls
- 16205963 LDoms-PM Obs Mod needs to close sockets in the error paths
- 16209808 Obs Mod power polling code should support a new packed data format OID
- 16217494 PM does not handle padding of lesser state resources correctly
- 16219418 Warning: (7) Getting CPU chip state from the HV
- 16225577 Assertion failed: rio == 0 || rio->evacuated, file ldomcli/io\_cmds.c:ldm\_add\_bus
- 16230078 PM should report degraded resources differently than inaccessible resources
- 16237203 dr\_cpu\_unconfigure leaks buf from libds\_chan\_create\_rsp, client\_loop leaks conn
- 16237771 Status for empty EMS slot should be "EMP"
- 16248520 pm\_stop() free invalid or corrupted buffer
- 16291759 Logical Domains Manager returns success for certain types of destroy-vf failures on InfiniBand device
- 16292272 Domain service response timeouts and extreme slowness in LDoms operations on SPARC M4
- 16293078 ldm add-io operation fails with error messages but returns exit status 0
- 16297875 Assert failed during DR: ((ppchain).startp) != 0 in remap\_pa\_ppriority()
- 16299503 Memory leak in pm\_test-1.0\_lib.so : pmtest\_find\_cookies
- 16320538 ldm set-domain extended-mapin-space= <domain> doesn't free the LDC memory frag

- 16324997 Warm-Migration should fail if vdsdev in use by another ldom
- 16364884 ldm should reject fuzzed XML input
- 16371765 Extra tlb nodes left in GMDs bloat ldom config sizes for large systems
- 16388201 papsat calculation of % time in each coherency link state is broken
- 16389751 PM code cleanup 2: Splitting source files and organizing header files
- 16398345 Additional latency for bound ttfc in elastic mode should be more accurate
- 16416658 LDoms IPS pkg build broken if “solaris” isn't first publisher on build machine
- 16418555 Improve performance of binding and other LDoms operations
- 16424206 Blacklisted strands can be added back to a domain
- 16425026 ldmd aborted in ldm add-mem
- 16440060 add-vnet/set-vnet should check for duplicate mac address
- 16463734 Guest ldoms are created without the pm\_boot\_policy property
- 16464461 default setting of extended-mapin-space need to be set to “on”
- 16470783 Unbound resources are not power managed in performance or elastic policies
- 16476994 papsat: Add a new policy change event
- 16478484 VIO operations can fail silently
- 16487198 ldmd dumps core/libc.so.1's strlen() on dynamic add-io operations on qlc device
- 16499865 Unable to set-io for virtual function through XMPP interface
- 16516403 PM merge breaks platform-specific shared-object modules
- 16516698 Migration of domains with cpu-arch=native is broken in 3.1.0.0.6
- 16528682 Parfait 1.1p2 reports memory leaks in pmi\_add\_pwr\_state\_padding()
- 16531613 hvctl.h missing some error definitions
- 16536036 Guest domains remain after switching config and powercycle
- 16538249 Refactor cross-CPU migration code to mig\_hwcaps.c
- 16539932 segv on platforms that don't support power mgmt
- 16557729 Add support for memory-links L0s state to PM
- 16576598 Unsupported dr-pdom messages should return an error
- 16576627 Blacklist evacuation must use a simplified set of policies
- 16576784 The “evacuation pending” status indicator should be removed

- 16576813 Blacklisted memory can be added back to a domain
- 16596372 Recovery Mode: Degraded config recovered only 2GB of memory for guests
- 16596468 Recovery Mode: Degraded mode systems should emit a message for all ldm cmds
- 16597626 Recovery Mode: ldmd crashed while removing root complex
- 16610442 MDSET\_UPDATES\_REQUEST times out on SPARC M5-32
- 16627335 Service enable/disable must use service start/stop timeout
- 16636935 ldmd needs to handle specifying memory with Terabyte unit
- 16637100 Logical Domains Manager core dump at seq\_sync\_mem+0x94
- 16679127 Update XML interface for Recovery Mode
- 16682148 NPRD: confusing message when removing PCIE dev from bound NPRD
- 16695158 ldmd fails to start, dumps core in om\_get\_mib\_pwr()
- 16709873 ldm ls-io 3.1 has bad virtual function state after a dubious destroy-vf op s10u11
- 16711110 Recovery Mode: Domains with missing virtual functions are not recovered
- 16716929 Recovery Mode: Control domain recovery failure leaves it in delayed reconf mode
- 16717785 Recovery Mode: All available memory is not used while recovering guests
- 16731561 ldmpower -pr needs to show system averages
- 16739634 Logical Domains Manager must enable dynamic SRIOV feature based on presence of a specific FW prop
- 16741226 range\_overlap\_mblock\_list() walks the free list incorrectly
- 16747033 NPRD: ldmd dumped core on frag\_hvmd\_assign when creating virtual function in del-reconf mode
- 16764181 Physical functions are not shown in ldm ls-io after reboot primary
- 16769782 ldmd aborts on Fujitsu M10 system with HVctl\_st\_badmd
- 16769854 Port 16440230 to 3.1
- 16777219 Fatal error: Domain has more memory than HV !
- 16782687 PM: Timing of PAD CLAIM msg could cause coredump in elastic util event process
- 16789490 Infinite Loop in testpwr cap: Checking for state change
- 16791578 PM Util Low notification fails with multiple domain configuration

- 
- 16793506 ldm list-bindings primary gives a blank output for IO: PSEUDONYM column
  - 16815101 Core dump in Observability while traversing master domain list
  - 16823571 Add Fujitsu M10 system Migration Group Support
  - 16827688 Parfait 1.2 catches two file descriptor leaks
  - 16836999 Blacklist evacuation - core remap incorrectly uses a blacklisted core as target
  - 16842975 Blacklist evacuation - failed to blacklist the strands in a partial core
  - 16908607 ldmd core dump at initiate\_or\_update\_delayed\_reconfig
  - 16939693 Need to mitigate HV Group MMU Demap bug for live migration
  - 16942468 ldmd can save a config with guest\_no\_reset set to true
  - 16957739 Logical Domains Manager rennumbers pcie\_device nodes
  - 16960497 ASSERT(size\_so\_far < req\_size); failed in affinity\_choose\_bindings() line 1081
  - 16973290 Add Oracle VM Server for SPARC 3.1 man pages
  - 16980537 Slot removal operations on bound domains fail with incorrect message
  - 16993217 SPARC M5 guest hang during reboot after mem row evacuation
  - 17027893 MAC address check with domains having 100's of MAC IDs cause ldom migration fail
  - 17052248 Expand add-vdsdev force option to mpgroup
  - 17179054 LDoms performance policy ttfc is causing bad performance on SPARC T5
  - 17215630 Support for Fujitsu Board DR
  - 17335156 Integrate Fujitsu Board DR Support
  - 17511365 Promote Fujitsu Board DR development feature to product feature

