

Oracle® VM Server for SPARC 2.2 Release Notes

Copyright © 2007, 2012, 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 pas conçu ni 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

Preface	5
1 Oracle VM Server for SPARC 2.2 Release Notes	9
What's New in This Release	9
System Requirements	10
Supported Platforms	10
Required Software and Patches	12
Related Software	18
Optional Software	18
Software That Can Be Used With the Logical Domains Manager	19
System Controller Software That Interacts With Logical Domains Software	20
Upgrading to the Oracle VM Server for SPARC 2.2 Software	20
Known Issues	21
General Issues	21
Domain Migration Restrictions	31
Oracle VM Server for SPARC MIB Issues	34
Bugs Affecting the Oracle VM Server for SPARC 2.2 Software	34
Resolved Issues	70

Preface

The *Oracle VM Server for SPARC 2.2 Release Notes* includes information about this release of the software, such as changes for this release, supported platforms, a matrix of required software and patches, and bugs affecting the Oracle VM Server for SPARC 2.2 software.

Related Documentation

The following table shows the documentation that is available for the Oracle VM Server for SPARC 2.2 release. These documents are available in HTML and PDF formats unless indicated.

TABLE P-1 Related Documentation

Application	Title
Oracle VM Server for SPARC 2.2 Software	<i>Oracle VM Server for SPARC 2.2 Administration Guide</i>
	<i>Oracle VM Server for SPARC 3.0 Security Guide</i>
	<i>Oracle VM Server for SPARC 2.2 Reference Manual</i>
	<i>Oracle VM Server for SPARC 2.2 Release Notes</i>
Oracle VM Server for SPARC 2.2 drd(1M) and vntsd(1M) man pages	Oracle Solaris OS Reference Manuals:
	<ul style="list-style-type: none">■ Oracle Solaris 10 Documentation■ Oracle Solaris 11 Documentation
Oracle Solaris OS: Installation and Configuration	Oracle Solaris OS Installation and Configuration Guides:
	<ul style="list-style-type: none">■ Oracle Solaris 10 Documentation■ Oracle Solaris 11 Documentation
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">■ <i>Secure Deployment of Oracle VM Server for SPARC</i> (http://www.oracle.com/technetwork/articles/systems-hardware-architecture/secure-ovm-sparc-deployment-294062.pdf)■ <i>Oracle Solaris 10 Security Guidelines</i>■ <i>Oracle Solaris 11 Security Guidelines</i>

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.

You can access the Oracle VM Server for SPARC discussion forum at <http://forums.oracle.com/forums/forum.jspa?forumID=1047>.

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.

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-2 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-3 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

Oracle VM Server for SPARC 2.2 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 2.2 software.

Important Installation Step – Prior to installation, you *must* unbind any domains that have the whole-core constraint set. After you complete the installation, rebind the domain. Optionally, you can save a service processor (SP) configuration. By unbinding and rebinding, the whole-core constraint is preserved.

Note – Oracle VM Server for SPARC features are added and maintained on the supported hardware platforms listed in [“Supported Platforms” on page 10](#). 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 T-Series SPARC servers at the time that the Oracle VM Server for SPARC releases and not for SPARC systems that have already passed their last-order date.

What's New in This Release

The major changes for this release of the Oracle VM Server for SPARC 2.2 software are as follows:

- Support for PCIe Single Root I/O Virtualization (SR-IOV) physical functions and virtual functions. See [“Using PCIe SR-IOV Virtual Functions” in *Oracle VM Server for SPARC 2.2 Administration Guide*](#).
- Support for performing named CPU core and memory block assignments. See [“Assigning Physical Resources to Domains” in *Oracle VM Server for SPARC 2.2 Administration Guide*](#).
- CPU whole-core allocation enhancement. See [“Configuring the System With Hard Partitions” in *Oracle VM Server for SPARC 2.2 Administration Guide*](#).

- Single-CPU performance enhancements that provide high instructions per cycle (IPC) on Oracle's SPARC T4 platform. See “[Tuning the SPARC CPU to Optimize Workload Performance on SPARC T4 Systems](#)” in *Oracle VM Server for SPARC 2.2 Administration Guide*.
- Cross-CPU migration support. For information about the `cpu-arch` property, see the `ldm(1M)` man page.
- Virtual I/O support for Rx Dring mode. For information about the `extended-mapin-space` property, see the `ldm(1M)` man page.
- Bug fixes

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>).

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>).

Note – Starting with Logical Domains 1.3, UltraSPARC T1 platforms are no longer supported. Earlier releases of the Logical Domains software continue to support these platforms. By default, the Oracle Solaris 11 OS includes Version 2.1 of the Logical Domains Manager.

However, Logical Domains Manager 2.2 is *not* supported on UltraSPARC T1 systems.

The Oracle VM Server for SPARC 2.2 software is supported on the following platforms:

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

Required Software and Patches

This section lists the required software and patches for use with the Oracle VM Server for SPARC 2.2 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 and Recommended Oracle Solaris OS

To use all features of the Oracle VM Server for SPARC 2.2 software, the operating system (OS) on all domains should be at least the Oracle Solaris 10 8/11 OS or the Oracle Solaris 11 OS. This OS can be either an initial installation of or an upgrade to the Oracle Solaris 10 8/11 OS or the Oracle Solaris 11 OS.

Note – Some features, such as PCIe SR-IOV and cross-CPU migration (cpu-arch property), are only available with the Oracle Solaris 11 OS.

The minimum version of the Oracle Solaris 11 OS for each domain type is Oracle Solaris 11 Support Repository Update 7 (SRU 7). The following table shows the patches that are required for older releases of the Oracle Solaris 10 OS to have the Oracle VM Server for SPARC 2.2 functionality. These patches are included in the Oracle Solaris 10 8/11 OS.

TABLE 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
141514-02 (vntsd)	X	X		
144500-19 (Oracle Solaris 10 8/11 feature kernel update)	X	X	X	X
148233-02 (Oracle Solaris 10 8/11: SR-IOV and DIO only)	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 Oracle VM Server for SPARC 2.2 Features

To enable all the Oracle VM Server for SPARC 2.2 features, you must run at least version 7.4.2 of the system firmware on both UltraSPARC T2 servers and UltraSPARC T2 Plus servers. You must run at least version 8.2.0 of the system firmware on SPARC T3 servers. You must run at

least version 8.2.0 of the system firmware on SPARC T4 servers. This firmware is preinstalled on the SPARC T4 servers. For information about the required Oracle Solaris OS, see [“Required and Recommended Oracle Solaris OS”](#) on page 12.

Required and Recommended System Firmware Patches

To take advantage of all features of Oracle VM Server for SPARC 2.2, ensure that your server from Oracle runs at least these revisions of the following system firmware patches:

147307-04	Sun SPARC Enterprise T5120 and T5220 Servers
147308-03	Sun Blade T6320 Server Module
147309-04	Netra T5220 Server
147310-04	Sun SPARC Enterprise T5140 and T5240 Servers
147311-03	Sun SPARC Enterprise T5440 Server
147312-03	Sun Blade T6340 Server Module
147313-03	Netra T5440 Server
147314-03	Sun Netra T6340 Server Module
148816-01	SPARC T3-1 Server
148817-01	SPARC T3-2 Server
148818-01	SPARC T3-4 Server
148819-01	SPARC T3-1B Server
148820-01	Netra SPARC T3-1 Server
148821-01	Netra SPARC T3-1B Server
148822-01	SPARC T4-1 Server
148823-01	SPARC T4-2 Server
148824-01	SPARC T4-4 Server
148825-01	SPARC T4-1B Server
148826-01	Netra SPARC T4-1 Server
148827-01	Netra SPARC T4-2 Server
148828-01	Netra SPARC T4-2B Server

Minimum Version of Software Required

Running the Oracle VM Server for SPARC 2.2 software along with at least the minimum versions of the other software components specified in [Table 1–2](#) is supported. In such a configuration, you might not get all the features of the Oracle VM Server for SPARC 2.2 software. For production environments, it is best to run with the recommended system firmware version described in “[Required Software to Enable Oracle VM Server for SPARC 2.2 Features](#)” on [page 12](#), and with the Oracle Solaris 10 8/11 OS or the Oracle Solaris 11 OS.

An alternate OS strategy is to do one of the following:

- Upgrade the control and service domains to the Oracle Solaris 10 8/11 OS (or to Oracle Solaris 11 OS for SPARC T4 systems).
- Apply the patches listed in [Table 1–1](#).

You can continue running the guest domains at their existing patch level.

Following is a matrix of the minimum versions of required software. The Oracle VM Server for SPARC 2.2 package can be applied 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). See the data sheet for your platform at <http://www.oracle.com/technetwork/documentation/sparc-tseries-servers-252697.html>.

TABLE 1–2 Minimum Versions of System Firmware and OS Software

Supported Server	System Firmware	Oracle Solaris OS
SPARC T4 servers	8.1.1	See the data sheet for your SPARC T4 server
SPARC T3 servers	8.1.x	See the data sheet for your SPARC T3 server
UltraSPARC T2 Plus servers	7.4.x	See the data sheet for your UltraSPARC T2 Plus server
UltraSPARC T2 servers	7.4.x	See the data sheet for your UltraSPARC T2 server

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) feature to assign direct I/O devices to domains, you must run the appropriate software and use supported PCIe cards.

- 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 [Doc 1325454.1: Oracle VM Server for SPARC Direct I/O Feature \(https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=REFERENCE&id=1325454.1\)](https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=REFERENCE&id=1325454.1).

- Software Requirements.** To use the DIO feature, the following domains must run the supported OS:
 - primary domain.** At least the Oracle Solaris 10 9/10 OS plus patch ID 145868-01 or the Oracle Solaris 11 OS.
 - I/O domain.** Any Oracle Solaris OS that is supported by the platform.

PCIe SR-IOV Hardware and Software Requirements

Starting with the Oracle VM Server for SPARC 2.2 release, the PCIe SR-IOV feature is supported on SPARC T3 and SPARC T4 platforms.

- Hardware Requirements.** To use the SR-IOV feature, you can use on-board PCIe SR-IOV devices as well as PCIe SR-IOV plug-in cards. The supported on-board device is the Intel 1-Gbps Ethernet network device.

The following table shows the names of the PCIe network cards that can be used with the PCIe SR-IOV feature.

PCIe Card Name	Part Number	Supported Servers From Oracle
Sun Dual 10-Gigabit Ethernet SFP+ LP	(X)1109A-Z	SPARC T3-1, SPARC T3-2, SPARC T4-1, SPARC T4-2, Netra SPARC T3-1, Netra SPARC T4-1, Netra SPARC T4-2
Sun Dual 10-Gigabit SFP+ Ethernet Module	(X)1110A-Z	Sun Blade 6000, Sun Blade 6048, SPARC T3-4, SPARC T4-4, SPARC SuperCluster T4-4, Netra SPARC T3-1B, 6000 Modular System
Dual 10-Gigabit Ethernet Fabric Expansion Module	(X)4871A-Z	SPARC T3-1B, SPARC T4-1B, Netra SPARC T3-1B, Netra SPARC T4-1B

- Firmware Requirements.** To use the SR-IOV feature, the SPARC T3 and SPARC T4 systems must run at least version 8.2.0 of the 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.

- **Software Requirements.** To use the SR-IOV feature, the following domains must run the supported Oracle Solaris OS:
 - The primary domain must be running the Oracle Solaris 11 OS (at least Support Repository Update 7 (SRU 7)).
 - The I/O domains can run the Oracle Solaris 10 9/10 OS, Oracle Solaris 10 8/11 OS, or Oracle Solaris 11 OS with at least SRU 7. For the Oracle Solaris 10 9/10 and Oracle Solaris 10 8/11 releases, install the `igbvf` and `ixgbev` virtual function driver patches, which are 148225 and 148227, respectively.

Live Migration Requirements

The *live migration* feature provides performance improvements that enable an active domain to be migrated while it continues to run. See [Chapter 9, “Migrating Domains,” in Oracle VM Server for SPARC 2.2 Administration Guide](#).

To take advantage of these migration performance improvements, the source and target machines *must* run at least these versions of the software components:

- **Logical Domains Manager:** At least Oracle VM Server for SPARC 2.1
- **System firmware:**
 - Version 7.4.0 for UltraSPARC T2 and UltraSPARC T2 Plus systems
 - Version 8.1.0 for SPARC T3 systems

- At least version 8.1.0 for SPARC T4 systems

In addition, the domain being migrated *must* run at least the Oracle Solaris 10 9/10 OS or the Oracle Solaris 11 OS.

Location of the Oracle VM Server for SPARC 2.2 Software

You can obtain the latest packages for both the Oracle Solaris 10 OS and Oracle Solaris 11 OS. 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-2_2.zip` package from My Oracle Support. See “How to Download the Logical Domains Manager Software (Oracle Solaris 10)” in *Oracle VM Server for SPARC 2.2 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 2.2 Software (Oracle Solaris 11)” in *Oracle VM Server for SPARC 2.2 Administration Guide*. Also see 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>).

For Oracle Solaris 10, the `OVM_Server_SPARC-2_2.zip` file that you download contains the following:

- Oracle VM Server for SPARC 2.2 software (`SUNWldm.v`)
- The `ldm(1M)`, `ldmconfig(1M)`, and `ldmd(1M)` man pages in the `SUNWldm.v` package are installed when the package is installed
- Installation script for Oracle VM Server for SPARC 2.2 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:

```
OVM_Server_SPARC-2_2/
  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

The *Oracle VM Server for SPARC 2.2 Administration Guide*, the *Oracle VM Server for SPARC 2.2 Reference Manual*, and these *Oracle VM Server for SPARC 2.2 Release Notes* can be obtained from:

<http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html>

For the Oracle Solaris 10 OS, the Oracle VM Server for SPARC 2.2 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 2.2 man pages, and the Japanese translations, are installed on your system as part of the ldomsmanager package.

For information about the Oracle Solaris OS, see the following libraries on the Oracle Technology Network:

- [Oracle Solaris 10 Documentation \(http://www.oracle.com/technetwork/documentation/solaris-10-192992.html\)](http://www.oracle.com/technetwork/documentation/solaris-10-192992.html)
- [Oracle Solaris 11 Documentation \(http://www.oracle.com/technetwork/documentation/solaris-11-192991.html\)](http://www.oracle.com/technetwork/documentation/solaris-11-192991.html)

Related Software

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 15, “Using the Oracle VM Server for SPARC Management Information Base Software,” in *Oracle VM Server for SPARC 2.2 Administration Guide*.

Software That Can Be Used With the Logical Domains Manager

This section details the software that is compatible with and can be used with the Logical Domains 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 Logical Domains software and your platform.

- **SunVTS** functionality is available in the control domain and guest domains on certain Logical Domains software 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 the *SunVTS 7.0 Software*.
- **Explorer Data Collector** can be used with the Logical Domains Manager 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 the *Oracle Explorer User's Guide*.
- **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. Starting with the Logical Domains 1.2 release and the Oracle Solaris Cluster 11/09 release, you can also manage logical domains as a resource by using the Oracle Solaris Cluster failover agent.
- **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 Containers and Logical Domains, and supports hardware management from power up to production. For more information, see <http://www.oracle.com/us/products/enterprise-manager/opscenter/index.html>.

System Controller Software That Interacts With Logical Domains Software

The following system controller (SC) software interacts with the Oracle VM Server for SPARC 2.2 software:

- **Sun Integrated Lights Out Manager (ILOM) 3.0** is the system management firmware that you can use to monitor, manage, and configure UltraSPARC T2, UltraSPARC T2 Plus, SPARC T3, and SPARC T4 server platforms. ILOM is preinstalled on these platforms and can be used on supported servers with the Oracle VM Server for SPARC 2.2 software enabled. Refer to the *Sun Integrated Lights Out Manager 3.0 User's Guide* for features and tasks that are common to Oracle's Sun rackmounted 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. Refer to the *Netra Data Plane Software Suite 2.0 User's Guide* for more information about this software.

Upgrading to the Oracle VM Server for SPARC 2.2 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.

Starting with the Logical Domains 1.0.1 release, you can upgrade to the Oracle VM Server for SPARC 2.2 software, see “[How to Upgrade to the Oracle VM Server for SPARC 2.2 Software \(Oracle Solaris 10\)](#)” in *Oracle VM Server for SPARC 2.2 Administration Guide* or “[How to Upgrade to the Oracle VM Server for SPARC 2.2 Software \(Oracle Solaris 11\)](#)” in *Oracle VM Server for SPARC 2.2 Administration Guide*.

Note – Starting with Logical Domains 1.3, UltraSPARC T1 platforms are no longer supported. Earlier releases of the Logical Domains software continue to support these platforms. By default, the Oracle Solaris 11 OS includes version 2.2 of the Logical Domains Manager.

However, Logical Domains Manager 2.2 is *not* supported on UltraSPARC T1 systems.

Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 2.2 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.

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 2.2 Administration Guide*.

I/O MMU Bypass Mode Is No Longer Needed

Starting with the Oracle VM Server for SPARC 2.0 release, I/O memory management unit (MMU) bypass mode is no longer needed. As a result, the `bypass=on` property is no longer available for use by the `ldm add -io` command.

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 Oracle 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 8/11 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 8/11, 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, this can cause the Oracle Solaris Volume Manager to be unable to find its configuration or to access its metadevices.

Workaround: After upgrading a service domain to Oracle Solaris 10 8/11, if a guest domain is unable to find its Oracle Solaris Volume Manager configuration or its metadevices, execute the following procedure.

▼ Find a Guest Domain's Oracle Solaris Volume Manager Configuration or Metadevices

1 Boot the guest domain.

2 Disable the `dev: id` feature of Oracle Solaris Volume Manager by adding the following lines to the `/kernel/dr/md.conf` file:

```
md_dev: id_destroy=1;  
md_kee: repl_state=1;
```

3 Reboot the guest domain.

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

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

5 Re-enable the Oracle Solaris Volume Manager `dev: id` 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 dev: id for 'vdc', 0x10
```

These messages are normal and do not report any problems.

Logical Domain Channels and Logical Domains

There is a limit to the number of logical domain channels (LDCs) that are available in any logical domain. For UltraSPARC T2 servers, SPARC T3-1 servers, SPARC T3-1B servers, SPARC T4-1 servers, and SPARC T4-1B servers, the limit is 512. For UltraSPARC T2 Plus servers, the other SPARC T3 servers and the other SPARC T4 servers, the limit is 768. This only becomes an issue on the control domain because the control domain has at least part, if not all, of the I/O subsystem allocated to it. This might also be an issue because of the potentially large number of LDCs that are created for both virtual I/O data communications and the Logical Domains Manager control of the other logical domains.

If you try to add a service, or bind a domain, so that the number of LDC channels exceeds the limit on the control domain, the operation fails with an error message similar to the following:

13 additional LDCs are required on guest primary to meet this request, but only 9 LDCs are available

If you have a large number of virtual network devices that are connected to the same virtual switch, you can reduce the number of LDC channels assigned by using the `ldm add-vsw` or `ldm set-vsw` command to set `inter-vnet-link=off`. When this property is set to `off`, LDC channels are not used for inter-vnet communications. Instead, an LDC channel is assigned only for communication between virtual network devices and virtual switch devices. See the [ldm\(1M\)](#) man page.

Note – Although disabling the assignment of inter-vnet channels reduces the number of LDCs, it might negatively affect guest-to-guest network performance.

The following guidelines can help prevent creating a configuration that could overflow the LDC capabilities of the control domain:

1. The control domain allocates approximately 15 LDCs for various communication purposes with the hypervisor, Fault Management Architecture (FMA), and the system controller (SC), independent of the number of other logical domains configured. The exact number of LDC channels that is allocated by the control domain depends on the platform and on the version of the software that is used.
2. The control domain allocates 1 LDC to every logical domain, including itself, for control traffic.
3. Each virtual I/O service on the control domain consumes 1 LDC for every connected client of that service.

For example, consider a control domain and 8 additional logical domains. Each logical domain needs the following at a minimum:

- Virtual network
- Virtual disk
- Virtual console

Applying the above guidelines yields the following results (numbers in parentheses correspond to the preceding guideline number from which the value was derived):

$$15(1) + 9(2) + 8 \times 3(3) = 48 \text{ LDCs in total}$$

Now consider the case where there are 45 domains instead of 8, and each domain includes 5 virtual disks, 5 virtual networks, and a virtual console. Now the equation becomes:

$$15 + 46 + 45 \times 11 = 556 \text{ LDCs in total}$$

Depending upon the number of supported LDCs of your platform, the Logical Domains Manager will either accept or reject the configurations.

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 less 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](#).

The OpenBoot PROM has a minimum size restriction for a domain. Currently, that restriction is 12 Mbytes. If you have a domain less than that size, the Logical Domains Manager will automatically boost the size of the domain to 12 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 2.2 Administration Guide](#).

Booting a Large Number of Domains

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

- 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 over 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. Do not have more than 32 vnet instances per vsw service because having more than that tied to a single vsw could cause hard hangs in the service domain.

To run the maximum configurations, a machine needs the 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](#), and [Installing Oracle Solaris 11 Systems](#).

Memory and swap space usage increases in a guest domain when the vsw services used by the domain provides services to many virtual networks (in multiple domains). This is due to the

peer-to-peer links between all the vnet connected to the vsw. The service domain benefits from having extra memory. Four Gbytes is the recommended minimum when running more than 64 domains. Start domains in groups of 10 or less 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 channels. See [“Inter-Vnet LDC Channels”](#) in *Oracle VM Server for SPARC 2.2 Administration Guide*.

Cleanly Shutting Down and Power Cycling a Logical Domains 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 a Logical Domains system, make sure that you save the latest configuration that you want to keep.

▼ 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.

▼ 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 Logical Domains 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. This can be seen in the following example 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 powercycle, unless the variable updates are either initiated from OpenBoot firmware on the control domain or followed by saving the configuration to the SC.

In this context, it is important to note that a reboot of the control domain could initiate a powercycle of the system:

- When the control domain reboots, if there are no bound guest domains, and no delayed reconfiguration in progress, the SC powercycles the system.
- When the control domain reboots, if there are guest domains bound or active (or the control domain is in the middle of a delayed reconfiguration), the SC does not powercycle 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 `eeprom(1M)` command
- Using the Logical Domains Manager CLI (`ldm`)
- Modifying, in a limited fashion, from the system controller (SC) using the `bootmode` command, that is, only certain variables, and only when in the `factory-default` configuration

The goal is that, variable updates that are made by using any of these methods always persist across reboots of the domain. The variable updates also always reflect in any subsequent logical domain configurations that were saved to the SC.

In Oracle VM Server for SPARC 2.2 software, there are a few cases where variable updates do not persist as expected:

- All methods of updating a variable persist across reboots of that domain. However, they do not persist across a powercycle of the system, unless a subsequent logical domain configuration is saved to the SC. The methods of updating a variable include by OpenBoot firmware and by the `eeprom` and `ldm` commands. In addition, in the control domain, updates made using OpenBoot firmware persist across a powercycle of the system, that is, even without subsequently saving a new logical domain configuration to the SC.
- 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 `ntptime` 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: **6520041**, **6540368**, **6540937**, and **6590259**.

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.

Containers, Processor Sets, and Pools Are Not Compatible With CPU Power Management

Using CPU dynamic reconfiguration (DR) to power down virtual CPUs does not work with processor sets, resource pools, or the zone's dedicated CPU feature.

When using the CPU power management elastic policy, the Oracle Solaris OS guest sees only the CPUs that are allocated to the domains that are powered on. That means that output from the `psrinfo(1M)` command dynamically changes depending on the number of CPUs currently power-managed. This causes an issue with processor sets and pools, which require actual CPU IDs to be static to allow allocation to their sets. This can also impact the zone's dedicated CPU feature.

Workaround: Set the power management policy to the performance policy.

Fault Management

There are several issues associated with FMA and power-managing CPUs. If a CPU faults when running with the elastic policy set, switch to the performance policy until the faulted CPU recovers. If all faulted CPUs recover, then elastic policy can be used again.

Delayed Reconfiguration

When a primary domain is in a delayed reconfiguration state, CPUs are power managed only after the primary domain reboots. This means that CPU power management will not bring additional CPUs online while the domain is experiencing high-load usage until the primary domain reboots, clearing the delayed reconfiguration state.

Cryptographic Units

The Oracle Solaris 10 10/09 OS introduces the capability to dynamically add and remove cryptographic units from a domain, which is called cryptographic unit dynamic reconfiguration (DR). 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 are running an appropriate version of the Oracle Solaris OS.

No core disable operations are performed on domains that have cryptographic units bound when the SP is set to the elastic policy. To enable core disable operations to be performed when the system has the elastic policy set, remove the cryptographic units that are bound to the domain.

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: normaal
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

```

Extended Mapin Space Is Only Available in the Oracle Solaris 10 8/11 OS and Oracle Solaris 11 OS

Extended mapin space is *only* available in the Oracle Solaris 10 8/11 OS and Oracle Solaris 11 OS. By default, this feature is disabled.

You can use the `ldm add-domain` or `ldm set-domain` command to enable the mode by setting `extended-mapin-space=on` on a domain that is running the Oracle Solaris 10 8/11 OS or Oracle Solaris 11 OS. See the [ldm\(1M\)](#) man page.

Graphical Configuration Assistant Tool Has Been Removed

Starting with the Oracle VM Server for SPARC 2.1 release, only the terminal-based Configuration Assistant tool, `ldmconfig`, is available. The graphic user interface tool is no longer available.

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>
```

Block of Dynamically Added Memory Can Be Dynamically Removed Only as a Whole

A block of dynamically added memory can be dynamically removed only as a whole. That is, a subset of that memory block cannot be dynamically removed.

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

```

# ldm list ldom1
NAME STATE FLAGS  CONS VCPU MEMORY UTIL UPTIME
ldom1 active -n---- 5000 2    1G    0.4% 23h

# ldm add-mem 16G ldom1

```

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

```
# ldm rm-mem 16G ldom1
```

```
# ldm list ldom1
NAME STATE FLAGS CONS VCPU MEMORY UTIL UPTIME
ldom1 active -n---- 5000 2 1G 0.4% 23h
```

Workaround: Dynamically add memory in smaller amounts to reduce the probability that this condition will occur.

Recovery: Reboot the domain.

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 re-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.

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.

Version Restrictions for Migration

Both the source and target machines must run at least version 2.1 of the Logical Domains Manager.

The following examples show the messages that you see when you run older versions of the Logical Domains Manager, the system firmware, or both:

- The target machine runs an older version of the Logical Domains Manager.

For example, assume that the source and target machines are running the following:

- **Source machine.** Runs version 2.1 of the Logical Domains Manager and version 7.4 of the system firmware
- **Target machine.** Runs version 2.0 of the Logical Domains Manager and version 7.4 of the system firmware

```
# ldm migrate ldg1 system2
```

```
The target machine is running an older version of the domain
manager that does not support the latest migration functionality.
```

- The source machine runs an older version of the Logical Domains Manager.

For example, assume that the source and target machines are running the following:

- **Source machine.** Runs version 2.0 of the Logical Domains Manager and version 7.4 of the system firmware
- **Target machine.** Runs version 2.1 of the Logical Domains Manager and version 7.4 of the system firmware

```
# ldm migrate ldg1 system2
```

```
The source machine is running an older version of the domain manager
that is not compatible with the version running on the target machine.
```

- The source and target machines run an older version of the Logical Domains Manager.

For example, assume that the source and target machines are running the following:

- **Source machine.** Runs version 2.0 of the Logical Domains Manager and version 7.3 of the system firmware
- **Target machine.** Runs version 2.0 of the Logical Domains Manager and version 7.4 of the system firmware

```
# ldm migrate ldg1 system2
```

```
Unable to migrate guest resource state
Domain Migration of LDom ldg1 failed
```

- The target machine runs an older version of the system firmware that is not compatible with the version of the system firmware that runs on the source machine.

For example, assume that the source and target machines are running the following:

- **Source machine.** Runs version 2.1 of the Logical Domains Manager and version 7.4 of the system firmware
- **Target machine.** Runs version 2.1 of the Logical Domains Manager and version 7.3 of the system firmware

```
# ldm migrate ldg1 system2
```

The target machine is running an older version of the System Firmware that is not compatible with the version running on the source machine.

- The source machine runs an older version of the system firmware that is not compatible with the version of the system firmware that runs on the target machine.

For example, assume that the source and target machines are running the following:

- **Source machine.** Runs version 2.1 of the Logical Domains Manager and version 7.3 of the system firmware
- **Target machine.** Runs version 2.1 of the Logical Domains Manager and version 7.4 of the system firmware

```
# ldm migrate ldg1 system2
```

The source machine is running an older version of the System Firmware that does not support the latest migration functionality.

CPU Restrictions for Migration

If the domain to be migrated is running an Oracle Solaris OS version older than the Oracle Solaris 10 8/11 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 *only* when you run an OS that precedes the Oracle Solaris 10 8/11 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 entirely 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 2.2 Administration Guide*.

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.

Note – The Oracle VM Server for SPARC MIB software is *only* available on Oracle Solaris 10 systems.

The `snmptable` Command Does Not Work With the Version 2 or Version 3 Option

Bug ID 6521530: 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 “[How to Retrieve Oracle VM Server for SPARC MIB Objects](#)” in *Oracle VM Server for SPARC 2.2 Administration Guide*.

Bugs Affecting the Oracle VM Server for SPARC 2.2 Software

This section summarizes the bugs that you might encounter when using this version of the software. The bug descriptions are in numerical order by bug ID. If a workaround and a recovery procedure are available, they are specified.

PCIe Fabrics Are Not Accessible by Guest Domains When 11 or More Domains Have PCIe Devices

Bug ID 7166620: If the control domain is rebooted when 11 or more guest domains have PCIe endpoint devices assigned, the PCIe devices are inaccessible on the guest domain.

Recovery: Stop and restart the affected guest domains.

Workaround: Configure a domain dependency relationship between the control domain and the guest domains that have PCIe endpoint devices assigned to them. The following dependency relationship ensures that domains that have PCIe endpoint devices are automatically stopped when the control domain reboots for any reason:

```
primary# ldm set-domain failure-policy=stop primary
primary# ldm set-domain master=primary ldom
```

Ldmd Terminates Abnormally on Operations After Canceling a Delayed Reconfiguration

Bug ID 7165095 and 7165101: On a system that has Direct I/O or SR-IOV domains, canceling a delayed reconfiguration and then performing any subsequent reconfiguration operation, the ldmd daemon terminates abnormally and produces a core file. The ldmd SMF service might also enter maintenance mode.

Workaround: Avoid the use of the `ldm cancel -reconf` command. If you must cancel or have already canceled the delayed reconfiguration, restart the ldmd SMF service before you perform any other ldm operations.

```
# svcadm restart ldmd
```

Recovery: If the ldmd SMF service enters maintenance mode, you must power cycle the system before you can restore the ldmd service.

The following show how to power cycle the system from the control domain and from the service processor (SP):

- **Control domain.** Perform the following command:

```
# halt
```

- **SP.** Perform the following commands:

```
-> stop /SYS
```

```
Are you sure you want to stop /SYS (y/n)? y
```

```
-> show /HOST status
```

```
/HOST
```

```
Properties:
```

```
status = Powered Off
```

```
-> start /SYS
```

```
Are you sure you want to start /SYS (y/n)? y
```

```
Starting /SYS
```

```
-> start /HOST/console
```

```
Are you sure you want to start /HOST/console (y/n)? y
```

```
->
```

Unbound Domain That Has Disabled CPUs Reports Incorrect Number of CPU Resources

Bug ID 7160502: Disabled CPUs can make the Logical Domains Manager report an incorrect number of CPU resources. The following example shows that unbinding a domain erroneously changes the CPU resource count for the domain:

```
# ldm list
NAME          STATE      FLAGS  CONS  VCPU  MEMORY  UTIL  UPTIME
primary       active    -n-cv-  UART   9     4G      0.2%  1h 5m
ldg1          bound     - - - - - 5000   116    2G
# ldm unbind ldg1
# ldm list
NAME          STATE      FLAGS  CONS  VCPU  MEMORY  UTIL  UPTIME
primary       active    -n-cv-  UART   9     4G      1.1%  1h 5m
ldg1          inactive  - - - - -          120    2G
```

At this point, the number of CPU resources is incorrect. The count for the ldg1 domain should be 116 and not 120 as shown after the unbind operation.

Note – This example is just one example, and there might possibly be other situations where the CPU count is incorrect due to disabled CPUs. In such cases, use the approach that is presented in the workaround.

Workaround: If possible, avoid using cores that have disabled CPUs. Otherwise, when unbinding a domain that has disabled cores, take care to reset the number of CPUs to the correct amount so that the domain can later be rebound.

To rebound the domain, you must reset the number of CPU resources. For example:

```
# ldm set-vcpu 116 ldg1
# ldm bind ldg1
```

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

Bug ID 7159359: 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: Use the `ldm list-io` command to save the information about the virtual functions, and then use the `ldm rm-dom` command to destroy each affected domain. Then, use the `ldm create-vf` command to create all the required virtual functions. Now, you can use the `ldm` command to rebuild the domains. 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 Cannot Correctly Re-create a Domain That Has Virtual Function Devices” on page 37](#).

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

Bug ID 7159114: 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 Cannot Correctly Re-create a Domain That Has Virtual Function Devices

Bug ID 7158496: When you use the `ldm list-constraints -x` command to save constraints to an XML file, the virtual function details are not saved. As a result, when the configuration is reset to factory-default, and the `ldm init-system` command is run to re-create the saved configuration, the virtual functions are not created, and any domain bind attempts fail.

Workaround: If any existing configuration has virtual functions, save all information about those virtual functions. You can later use this information to manually re-create the virtual functions before running the `ldm init-system` command.

The following procedure shows how to save all the information about the virtual functions to use later:

1. Save the domain configuration in a file, `vfs.txt`, for use in re-creating the virtual functions.

```
primary# ldm list-io -l -p | grep "type=VF" >vfs.txt
```

A typical virtual function entry in `vfs.txt` appears as follows:

```
|dev=pci@400/pci@1/pci@0/pci@4/network@0,83|alias=/SYS/MB/NET0/IOVNET.PF1.VF1|
  status=RDY|domain=ldg1|type=VF|class=NETWORK
|proptype=class|mac-addr=00:14:4f:f9:74:d0
|proptype=class|vlan-ids=3,5,7
|proptype=class|mtu=1500
|proptype=device|unicast-slots=6
```

The first line is intentionally split into two lines for readability. It will be a single line in the `vfs.txt` file.

2. Reset the domain to the factory-default configuration.
3. Reboot the control domain.
4. Create the virtual functions based on the information in the `vfs.txt` file.

For each such entry, use the `ldm create-vf` command to re-create the virtual function with its original name and properties. Use the following command for the example virtual function:

```
primary# ldm create-vf mac-addr=00:14:4f:f9:74:d0 vid=3,5,7 mtu=1500 \
unicast-slots=6 /SYS/MB/NET0/IOVNET.PF1
```

For details about the `class` and `device` properties, see the [ldm\(1M\)](#) man page.

Note – The virtual function name is generated from the name of its parent physical function. As a result, execute the `ldm create-vf` commands in increasing numeric order based on the virtual function part of the name. For example, physical function `/SYS/MB/NET0/IOVNET.PF1` has the following child virtual functions:

```
/SYS/MB/NET0/IOVNET.PF1.VF0 mac-addr=00:14:4f:f9:74:d0
/SYS/MB/NET0/IOVNET.PF1.VF1 mac-addr=00:14:4f:f9:74:d1
```

The following commands create the virtual functions:

```
primary# ldm create-vf mac-addr=00:14:4f:f9:74:d0 /SYS/MB/NET0/IOVNET.PF1
Created new VF: /SYS/MB/NET0/IOVNET.PF1.VF0
primary# ldm create-vf mac-addr=00:14:4f:f9:74:d1 /SYS/MB/NET0/IOVNET.PF1
Created new VF: /SYS/MB/NET0/IOVNET.PF1.VF1
```

The first `ldm create-vf` command causes the system to enter delayed reconfiguration mode.

5. Verify that the new configuration includes the virtual functions that you manually created.

```
primary# ldm list-io -l -p | grep "type=VF" >vfs.after.txt
```

Compare the contents of the `vfs.after.txt` file with the `vfs.txt` file.

6. Reboot the control domain.
7. Reconfigure a domain from an XML file.

```
primary# ldm init-system -i file.xml
```

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

Bug ID 7158454: 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, if many virtual network devices in the same virtual switch exist across many domains. Typically, this issue is seen with around 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 `ldmd` 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 init-system Reports a disk server not found Error

Bug ID 7155386: When an XML file contains both control domain and guest domain configurations, the `ldm init-system` command first configures the guest domains and then configures the control domain. In a factory default configuration that does not have a virtual disk server configured, attempting to add a virtual disk server device to the guest domains might fail with the following error:

```
Disk Server xxx not found
```

This failure occurs if the specified virtual disk server should be provided by the control domain.

Setting Unicast Slots to a Number That Exceeds the Maximum Causes the Value to Be Reset to 0

Bug ID 7155349: Setting unicast slots to a number that exceeds the maximum limit fails with an appropriate error message. However, the number of unicast slots are incorrectly and silently reset to 0.

Workaround: Specify a value for the number of unicast slots that is within the range of supported values.

Attempt to Exceed the Maximum Number of Unicast Slots of ixgbe Physical Functions and Virtual Functions Does Not Fail

Bug ID 7155282: When you attempt to set more unicast slots of ixgbe physical functions and virtual functions than is permitted by the maximum limit, the command succeeds. Attempting to exceed this maximum limit should fail, but it doesn't.

Use the following command to identify the maximum number of unicast slots that is supported by the device:

```
# ldm list-io -d pf-name
```

Then, ensure that the total number of unicast slots given to each virtual function in that physical function does not exceed that maximum value.

Control Domain Requires the Lowest Core in the System

Bug ID 7153060: 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

VCPU
  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%
```

Ldmd Daemon Does Not Come Online

Bug ID 7151847: The Service Management Facility (SMF) service for the ldmd daemon does not come online if the Oracle VM Server for SPARC 2.2 software is installed on a control domain that runs the Oracle Solaris OS versions that are Oracle Solaris 10 10/09 or earlier. This situation occurs because an explicit SMF dependency on the svc:/ldoms/agents SMF service has been added.

Workaround: Install patch ID 142909-17, which adds support for the svc:/ldoms/agents SMF service, ldmd, upon which ldmd depends.

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

Bug ID 7150793: 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 complete 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 7150209: 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 7149951: 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 `cpurack` 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 7149365: 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 7149323: 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
```

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

Bug ID 7148394: 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.

`ldm init-system -r -i XML-file` Does Not Reboot the primary Domain

Bug ID 7146725: When you use the `ldm init-system` command to install a domain from an XML configuration, the primary domain fails to reboot even though the `-r` option is specified.

Workaround: Manually reboot the primary domain.

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

Bug ID 7146423: 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.

Version 8.2.0 of the System Firmware Contains a New Version of the scvar Database

Bug ID 7144314: version 8.2.0 of the system firmware contains a new version of the scvar database, which reverts to defaults after the installation completes.

Workaround: Make note of the running Oracle VM Server for SPARC configuration or any changed system diagnostic properties prior to installing the system firmware. Use the ILOM show command. For example:

```
-> show /HOST/domain/configs
```

After you install the firmware and prior to powering on the system, use the ILOM set command. For example:

```
-> set /HOST/bootmode config=config-name
```

At this time, the Oracle VM Server for SPARC configurations are preserved. However, you *must* select whether to boot a particular configuration or the factory-default configuration.

The following property values revert to default values after you install the firmware:

```
/HOST
  Properties:
  autorunonerror
  ioreconfigure

/HOST/bootmode
  Properties:
  config

/HOST/diag
  Properties:
  error_reset_level
  error_reset_verbosity
  hw_change_level
  hw_change_verbosity
  level
```

```
mode
power_on_level
power_on_verbosity
trigger
verbosity

/HOST/domain/control
Properties:
auto-boot
boot_guests

/HOST/tpm
Properties:
enable
activate
forceclear

/SYS
Properties:
keyswitch_state

/SP/powermgmt
Properties:
policy
```

panic: BAD TRAP: occurred in module "pcie" due to an illegal access to a user address

Bug ID 7142913: After you bind and start 15 guest domains, the primary domain panics and the following error message is issued:

```
panic: BAD TRAP: occurred in module "pcie" due to an illegal access to a user address
```

The domains are configured as follows:

- **Guest domain.** Owns the igb and ixgbe virtual function devices. Also has the master property set to primary.
- **primary domain.** Has the failure-policy property set to stop.

Control Domain Reconfigured From an XML File Fails to Remove I/O Devices Properly

Bug ID 7134203: Existing I/O devices are not properly removed from the control domain when the control domain is reconfigured from an XML file by using the `ldm init-system` command. This situation might cause a binding failure on the guest domain if the control domain still has the PCIe leaf node devices bound to the control domain.

An Invalid vdsdev Backend Is Seen as a Valid Path

Bug ID 7131596: If you specify an incorrect vdsdev backend to the `ldm add-vdsdev` command, the resulting error message identifies the backend as a valid path:

```
# ldm add-vdsdev /wrong/path/file disk1@primary-vds0
Path /wrong/path/file is valid but not accessible on service domain primary
```

Workaround: Verify and, if necessary, correct the specified path.

After Disabling the Whole-Core Constraint, the Constraint Reappears After a primary Domain Reboot

Bug ID 7130693: After the whole-core constraint is disabled, the constraint reappears after a primary domain reboot.

This problem only occurs in the following circumstances:

- The primary domain is in delayed reconfiguration mode.
- The number of virtual CPUs that you specify with the `ldm set -vcpu` command (without the `-c` option) matches the number of virtual CPUs that was used to set the whole-core constraint before the delayed reconfiguration was initiated.

Workaround: Disable the whole-core constraint by specifying a different number of virtual CPUs.

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

Bug ID 7129252: 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: Perform the following steps:

1. Remove the PCIe bus from the root domain.

```
primary# ldm rm-io pci_0 primary
Initiating a delayed reconfiguration operation on the primary domain.
All configuration changes for other domains are disabled until the primary
domain reboots, at which time the new configuration for the primary domain
will also take effect.
```

2. Reassign the PCIe bus to the root domain.

```
primary# ldm add-io pci_0 primary
-----
Notice: The primary domain is in the process of a delayed reconfiguration.
Any changes made to the primary domain will only take effect after it reboots.
-----
```

3. Reboot the PCIe bus to the root domain.

primary# **reboot**

Ldm start Erroneously Returns 0 Instead of 1 on Failure to Start a Guest Domain

Bug ID 7125579: A guest domain might fail to start because of an unexpected hypervisor error. Even if the domain fails to start, the command exits with 0 instead of 1 and issues the following error message:

```
LDom domain start failed, retry the operation
```

Workaround: Do *not* rely solely on the exit code to determine whether a domain started successfully. Instead, perform one of the following checks:

- Verify that the `ldm` command issues an error message.
- Check the status of the domain after the start command completes.

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

Bug ID 7121963: Only use 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 re-add the card to the primary domain.

Ldm stop Command Might Fail If Issued Immediately After an Ldm start Command

Bug ID 7118936: 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.

Using Ldm set -io to Change pvid Value Twice in Succession Might Cause a Configuration Failure

Bug ID 7109458: Using the `ldm set -io` command to change the `pvid` property value for a virtual function more than one time might cause the `pvid` value to not be correctly set on the virtual function hardware.

Workaround: Wait a few seconds before you run the `ldm set -io` command again.

System Panics When Rebooting a primary Domain That Has a Very Large Number of Virtual Functions Assigned

Bug ID 7104911: A system panics 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 for all the `ixgbe` virtual functions on the system.

Vague SR-IOV Error Message: Create vf failed

Bug ID 7101229: When you attempt to create one more virtual function than the maximum number of configurable virtual functions for a physical function device, the `Create vf failed` message is issued. This error message is unclear as to the reason for the failure.

Oracle Solaris 11 OS: Using Direct I/O to Remove Multiple PCIe Slots From the primary Domain on a Multi-Socket SPARC T-Series System Might Panic at Boot Time

Bug ID 7100859: Your system might panic at boot time if you use direct I/O (`ldm remove-io`) to remove multiple PCIe slots from a multi-socket SPARC T-Series system. This occurs when the paths to the PCIe slots are similar to each other (except for the root complex path). The panic might occur after you remove the PCIe slots and then reboot the primary domain. For more information about the direct I/O (DIO) feature, see [“Assigning PCIe Endpoint Devices” in Oracle VM Server for SPARC 2.2 Administration Guide](#).

For example, if you remove the `/SYS/MB/PCIE5` (`pci@500/pci@2/pci@0/pci@0`) and `/SYS/MB/PCIE4` (`pci@400/pci@2/pci@0/pci@0`) slots, which have similar path names, the next boot of the Oracle Solaris 11 OS might panic.

The following `ldm list-io` command is run after the `/SYS/MB/PCIE4` and `/SYS/MB/PCIE5` PCIe slots are removed.

```
# ldm list-io
IO          PSEUDONYM      DOMAIN
--          -
pci@400     pci_0          primary
niu@480     niu_0          primary
pci@500     pci_1          primary
niu@580     niu_1          primary

PCIe          PSEUDONYM      STATUS  DOMAIN
-----          -

```

```

pci@400/pci@2/pci@0/pci@8 /SYS/MB/PCIE0 OCC primary
pci@400/pci@2/pci@0/pci@4 /SYS/MB/PCIE2 OCC primary
pci@400/pci@2/pci@0/pci@0 /SYS/MB/PCIE4 OCC
pci@400/pci@1/pci@0/pci@8 /SYS/MB/PCIE6 OCC primary
pci@400/pci@1/pci@0/pci@c /SYS/MB/PCIE8 OCC primary
pci@400/pci@2/pci@0/pci@e /SYS/MB/SASHBA OCC primary
pci@400/pci@1/pci@0/pci@4 /SYS/MB/NET0 OCC primary
pci@500/pci@2/pci@0/pci@a /SYS/MB/PCIE1 OCC primary
pci@500/pci@2/pci@0/pci@6 /SYS/MB/PCIE3 OCC primary
pci@500/pci@2/pci@0/pci@0 /SYS/MB/PCIE5 OCC
pci@500/pci@1/pci@0/pci@6 /SYS/MB/PCIE7 OCC primary
pci@500/pci@1/pci@0/pci@0 /SYS/MB/PCIE9 OCC primary
pci@500/pci@1/pci@0/pci@5 /SYS/MB/NET2 OCC primary
#

```

Workaround: Do *not* remove all slots that have similar path names. Instead, remove only one such PCIe slot.

You also might be able to insert the PCIe cards into slots that do not have similar paths and then use them with the DIO feature.

Partial Core primary Fails to Permit Whole-Core DR Transitions

Bug ID 7100841: 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.

After a primary Domain Reboot, igb and ixgbe Virtual Functions That Are Assigned to the primary Domain Become Faulty

Bug ID 7098941: The `igb` and `ixgbe` virtual function devices become faulty after the primary domain is rebooted. These virtual functions are assigned to the primary domain. The system configuration only has a primary domain. No guest domains or I/O domains are configured.

The `fmadm faulty` command shows that each virtual function device is faulty. The `fmadm repair` command enables you to recover from the faults, but the faulty state returns each time you reboot the primary domain.

Workaround: Use the `fmadm repair` command to recover from the faults each time you reboot the primary domain.

Ldmconfig Is Only Supported on Oracle Solaris 10 Systems

Bug ID 7093344: You can *only* use the `ldmconfig` command on Oracle Solaris 10 systems.

Ldm list -io Command Shows the UNK or INV State After Boot

Bug ID 7084728: 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 reply from the Oracle Solaris OS.

This problem has only been reported 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.

Cannot Detach Network Interface Card Driver

Bug ID 7083321: The `nwam` daemon holds a reference count on the network interface card (NIC) device node, so the NIC driver cannot be detached.

Workaround: Do *not* use the Automatic network configuration profile. Instead, use the DefaultFixed network configuration profile.

Oracle VM Server for SPARC MIB Is Only Supported on Oracle Solaris 10 Systems

Bug ID 7082776: You can *only* use the Oracle VM Server for SPARC MIB on Oracle Solaris 10 systems.

Migrating a Very Large Memory Domain on SPARC T4-4s Results in a Panicked Domain on the Target System

Bug ID 7071426: 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, ffffffff, 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

Bug ID 7062298: 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.

A Large-Memory Domain in Elastic Mode Might Take a Long Time to Stop

Bug ID 7058261: When you use the `ldm stop` command to stop a large-memory domain while the system has the power management elastic policy set, it might take a long time. If the domain is sufficiently idle, the majority of the CPU threads that are assigned to the domain will be disabled. By disabling the CPUs, the processing that is required to stop a domain is left to the remaining active threads.

For example, a guest domain that has 252 Gbytes of memory and only 2 CPUs enabled takes approximately 7 minutes to stop.

Workaround: Disable power management (PM) by switching from the elastic policy to the performance policy before you stop the domain.

Cannot Use Oracle Solaris Hot Plug Operations to Hot Remove a PCIe Endpoint Device

Bug ID 7054326: You *cannot* use Oracle Solaris hotplug 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 2.2 Administration Guide](#).

Virtual Disk Validation Fails for a Physical Disk With No Slice 2

Bug ID 7042353: If a physical disk is configured with a slice 2 that has a size of 0, you might encounter the following issues:

- If you use the `ldm add-vdsdev` command to add such a disk as a back end for a virtual disk, the command fails:

```
# ldm add-vdsdev /dev/dsk/c3t1d0s2 vol@primary-vds0
Path /dev/dsk/c3t1d0s2 is not valid on service domain primary
```

You can work around this problem by installing patch ID 147708-01 on the primary domain and any service domains, and then restart the `svc:/ldoms/agents` service.

- If you use the `ldm bind` command to bind a domain that has such disk as a back end of a virtual disk, the command fails:

```
# ldm bind ldg3
Path /dev/dsk/c3t1d0s2 is not valid on service domain primary
```

You can work around this problem by using the `-q` option of the `ldm bind` command:

```
# ldm bind -q ldg3
```

Another workaround permits you to permanently disable the disk validation that is performed by the `ldm add -vdsdev` and `ldm bind` commands. As a result, you do not need to specify the `-q` option. Permanently disable the disk validation by updating the `device_validation` property of the `ldmd` service:

```
# svccfg -s ldmd setprop ldmd/device_validation=value
# svcadm refresh ldmd
# svcadm restart ldmd
```

Specify a value of `0` to disable validation for network and disk devices. Specify a value of `1` to disable validation for disk devices but still enable validation for network devices.

The possible values for the `device_validation` property are:

- `0` Disable validation for all devices
- `1` Enable validation for network devices
- `2` Enable validation for disk devices
- `3` Enable validation for network and disk devices
- `-1` Enable validation for all type of devices, which is the default

nxge Panics When Migrating a Guest Domain That Has Hybrid I/O and Virtual I/O Virtual Network Devices

Bug ID 7038650: 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 7036137: 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.

ldmd Fails to Remove Cores From a Domain That Has Partial Cores

Bug ID 7035438: `ldmd` permits you to enable the whole-core constraint on a domain that has partial cores, yet fails to remove or set cores from the same domain.

Workaround: Do the following from the factory-default configuration on the control domain:

1. Initiate a delayed reconfiguration on the control domain.

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

2. Perform any memory reconfiguration operations first.

3. Perform the CPU reconfiguration operations.

```
# ldm set-vcpu 16 primary
# ldm set-vcpu -c 2 primary
```

This example uses 2 cores but the number of cores can be from 1 to the system limit.

Logical Domains Agent Service Does Not Come Online if the System Log Service Does Not Come Online

Bug ID 7034191: 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 Hang During a Migration

Bug ID 7030045: 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 powercycle 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 7027105: 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 there are more CPU resources allocated to the domain than is shown in the `psrinfo` 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 `psrinfo` output.

Live Migration of a Domain That Depends on an Inactive Master Domain on the Target Machine Causes `ldmd` to Fault With a Segmentation Fault

Bug ID 7026177: 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. You can still perform a migration, but 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 7026160: You perform a domain migration while a DRM policy is in effect. Later, 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 7025445: 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 7023216: 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 7019493: If two `ldm migrate` commands are issued simultaneously in the “opposite direction,” the two commands might hang and never complete. For example, an opposite direction situation is one where you simultaneously start a migration on machine A to machine B and a migration on machine B to machine A.

The hang results for the migration processes even if they are initiated as dry runs by using the `-n`. When this problem occurs, all other `ldm` commands might hang.

Workaround: None.

Removing a Large Number of CPUs From the Control Domain

Bug ID 6994984: 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 31](#)). Instead, undo the delayed reconfiguration operation by using the `ldm cancel-reconf` command, and start over.

3. Reboot the control domain.

System That Has the Elastic Policy Set and Is Running the Oracle Solaris 10 8/11 OS Might Hang

Bug IDs 6989192 and 7071760: 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 6984681: 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: Detect And Handle Disks That Are Accessible Through Multiple Direct I/O Paths

Bug ID 6984008: 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 impact 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. The steps to determine if you have dual-ported disks on T3-1 system are as follows:

Determine whether the system has dual-ported disks by running 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
```

When 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 6983279: 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.

Recovery: Cancel the `ldm rm-mem` or `ldm set-mem` command.

Workaround: Cancel the `ldm rm-mem` or `ldm set-mem` command, and check if 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 `ldm stop -a` Command on Domains in a Master-Slave Relationship Leaves the Slave With the Stopping Flag Set

Bug ID 6979574: 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 6968507: 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 `attack` 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 6968100: 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 6967799: The `ldmconfig` script cannot properly create a stored logical domains 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 6965758: The migration of an active domain can fail if it is running a release older than the Oracle Solaris 10 10/09 OS 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 as 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 6956431: After an Oracle Solaris 10 9/10 domain has been suspended as part of a migration operation, memory dynamic reconfiguration (DR) is disabled. This applies 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 6936833: 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 6904849: 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 only contain 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 6904240: 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 0x80000000  
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 6897743: 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.

Then, all `ssh` connections will no longer use the hardware cryptographic units (and thus not benefit from the associated performance improvements), and `ssh` connections would not be disconnected when the cryptographic units are removed.

PCI Express Dual 10-Gigabit Ethernet Fiber Card Shows Four Subdevices in `ldm ls-io -l` Output

Bug ID 6892229: When you run the `ldm ls-io -l` command on a system that has an 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.

ldm Commands Are Slow to Respond When Several Domains Are Booting

Bug ID 6855079: 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 Netra T5440 Server.

Guest Domain Might Fail to Successfully Reboot When a System Is in Power Management Elastic Mode

Bug ID 6853273: While a system has the power management elastic policy set, rebooting a guest domain might produce the following warning messages and fail to boot successfully:

```
WARNING: /virtual-devices@100/channel-devices@200/disk@0:
Sending packet to LDC, status: -1
WARNING: /virtual-devices@100/channel-devices@200/disk@0:
Can't send vdisk read request!
WARNING: /virtual-devices@100/channel-devices@200/disk@0:
Timeout receiving packet from LDC ... retrying
```

Workaround: If you see these warnings, perform one of the workarounds in the following order:

- If the guest domain shows an `ok>` prompt and accepts input, type `reset -all`
- From the control domain, issue an `ldm stop domain-name` command, then issue an `ldm start domain-name` command
- Change the Power Management policy to the performance policy, stop and start the affected guest domain, and then return to the elastic policy

Guest Domain Sometimes Fails to Make Proper Domain Services Connection to the Control Domain

Bug ID 6839787: Sometimes, a guest domain that runs at least the Oracle Solaris 10 10/08 OS does not make a proper Domain Services connection to a control domain that runs the Oracle Solaris 10 5/09 OS.

Domain Services connections enable features such as dynamic reconfiguration (DR), FMA, and power management (PM). Such a failure occurs when the guest domain is booted, so rebooting the domain usually clears the problem.

Workaround: Reboot the guest domain.

Oracle Solaris 11: Zones Configured With an Automatic Network Interface Might Fail to Start

Bug ID 6837615: 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.

The following are the workarounds:

- **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 inter-zone 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 6836587: 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 should not occur during normal operation, the error was seen when the instance number of a virtual network device did 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 – Be aware of the warning contained in the man page that states “changes should not be made to `/etc/path_to_inst` without careful consideration.”

Newly Added NIU/XAUI Adapters Are Not Visible to Host OS If Logical Domains Is Configured

Bug ID 6829016: When Logical Domains is configured on a system and you add another XAUI network card, the card is not visible after the machine is powercycled.

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 6808832: 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+ 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 6781589: 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.

Constraint Database Is Not Synchronized to Saved Configuration

Bug ID 6773569: After switching from one configuration to another (using the `ldm set -config` command followed by a powercycle), domains defined in the previous configuration might still be present in the current configuration, in the inactive state.

This is a result of the Logical Domains Manager's constraint database not being kept in sync with the change in configuration. These inactive domains do not affect the running configuration and can be safely destroyed.

Migration Does Not Fail If a vdsdev on the Target Has a Different Back End

Bug ID 6772120: 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: Ensure that when you are configuring the target domain to receive a migrated domain 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:

- Do the following:
 - Migrate the domain back to the source machine.
 - Fix the vdsdev on the target to point to the correct disk back end.
 - 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:
 - Use the `ldm rm-vdisk` command to remove the disk.
 - Fix the vdsdev.
 - 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 6772089: In certain situations, a migration fails and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This 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 6764613: 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 6760933: 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 6757486: 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 an Logical Domains System Does Not Return the System to the OK Prompt

Bug ID 6753683: Sometimes, executing the `uadmin 1 0` command from the command line of an 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 6742805: 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 it takes to shut down a domain.

Workaround: Ensure that large memory configurations (>100 Gbytes) have at least one core. This results in a much faster shutdown time.

If the Oracle Solaris 10 5/08 OS Is Installed on a Service Domain, Attempting a Net Boot of the Oracle Solaris 10 8/07 OS on Any Guest Domain Serviced by It Can Hang the Installation

Bug ID 6705823: Attempting a net boot of the Oracle Solaris 10 8/07 OS on any guest domain serviced by a service domain running the Oracle Solaris 10 5/08 OS can result in a hang on the guest domain during the installation.

Workaround: Patch the miniroot of the Oracle Solaris 10 8/07 OS net install image with Patch ID 127111-05.

Simultaneous Net-Installation of Multiple Domains Fails When in a Common Console Group

Bug ID 6656033: 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.

The `scadm` Command Can Hang Following an SC or SP Reset

Bug ID 6629230: The `scadm` command on a control domain running at least the Solaris 10 11/06 OS can hang following an SC reset. The system is unable to properly reestablish a connection following an SC reset.

Workaround: Reboot the host to reestablish connection with the SC.

Recovery: Reboot the host to reestablish connection with the SC.

`ldc_close: (0xb) unregister failed, 11 Warning Messages`

Bug ID 6610702: You might see the following warning message on the system console or in the system log:

```
ldc_close: (0xb) unregister failed, 11
```

Note that the number in parentheses is the Oracle Solaris internal channel number, which might be different for each warning message.

Workaround: You can ignore these messages.

Guest Domain With Too Many Virtual Networks on the Same Network Using DHCP Can Become Unresponsive

Bug ID 6603974: 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.

Logical Domains Manager Does Not Retire Resources On Guest Domain After a Panic and Reboot

Bug ID 6591844: If a CPU or memory fault occurs, the affected domain might panic and reboot. If the Fault Management Architecture (FMA) attempts to retire the faulted component while the domain is rebooting, the Logical Domains Manager is not able to communicate with the domain, and the retire fails. In this case, the `fmadm faulty` command lists the resource as degraded.

Recovery: Wait for the domain to complete rebooting, and then force FMA to replay the fault event by restarting the fault manager daemon (`fmd`) on the control domain by using this command:

```
primary# svcadm restart fmd
```

OpenBoot PROM Variables Cannot be Modified by the `eeprom(1M)` Command When the Logical Domains Manager is Running

Bug ID 6540368: This issue is summarized in “[Logical Domains Variable Persistence](#)” on [page 26](#) and affects only the control domain.

Cannot Set Security Keys With Logical Domains Running

Bug ID 6510214: In a Logical Domains environment, there is no support for 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 only valid for a single use.

Behavior of the `ldm stop-domain` Command Can Be Confusing

Bug ID 6506494: There are some cases where 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
```

Resolved Issues

The following RFEs and bugs have been fixed for the Oracle VM Server for SPARC 2.2 software release:

- | | |
|---------|--|
| 6617215 | ldmd shows MAUs warning when starting up on system which does not have any MAUs |
| 6756315 | Domain migration needs to handle multipathed disks correctly |
| 6796886 | ldm should allow cpu pinning by id |
| 6912155 | Memory DR should support > 31 memory segments configuration |
| 6933033 | RFE: support named cores via CLI for inactive domains |
| 6936378 | RFE: Disable CPU defragmentation in LDOMs |
| 6963491 | Batch requests and provide both normalized and non-normalized cpu utilization to LDOMs PM from HV |
| 6964197 | RFE: Logical Domains Manager SR-IOV support |
| 6978991 | RFE: Add explicit CLI to set the whole core hard cap |
| 6994853 | Implement <code>affinity_reserve()</code> and <code>memory_allocate()</code> to decouple memory routines to allow for Affinity |
| 6999051 | Whole-core constraint lost after powercycle (no HV MD property) |
| 7006175 | PM test harness has illicit knowledge of <code>pm_cookie</code> |

- 7010284 RFE: Pwrcap enhancements for SPARC PAD
- 7011393 RFE: Add PAP SAT support to gather PPFE-PPSE pm state info for M4/T5 platforms
- 7011573 Support for low-CPI workloads (aka dynamic threading)
- 7011574 Binding a domain should use new cores if any are free; only share cores where that's not possible
- 7012639 Remove check of VID=0.
- 7016545 RFE: Link L1 memory pm support for M4/T5 platforms
- 7017748 Implement affinity_pm to support CPU defragmentation
- 7017760 Common routines for Frag and HVFrag
- 7018614 Remap around partial cores other than the lowest core for control domain to switch to whole-core
- 7022021 PRI domain service is lost after migrating back to a previous source machine
- 7023018 Add hvfrag support for new, per-IOS allocation class
- 7024499 Vds_Validation: ldmd failed to bind -f/-q guest using xml file with invalid vdsdev
- 7026362 avg_guest_util() and avg_cpu_util() confuse ns and ms
- 7029957 ldmp2v_convert should notice vntsd is in maintenance mode.
- 7030027 ldmconfig attempts to configure mau on T4 platform
- 7030070 PM does not manage vcpus in Elastic mode after a guest starts/reboots
- 7030098 MEMDR set-mem 8G primary fails under some memory configurations
- 7030601 get_state_md_template() leaks filename string
- 7030824 RFE: Support for migration between different CPU types and speeds
- 7031177 set-vsw second time core dump.
- 7033167 ldmd core dumps when attempting a migration when running on downrev firmware
- 7033252 RFE: LDDoms MIB - support for the new properties such as extended-mapin-space and low CPI
- 7033352 prtdiag shows dual cpu speed after a few migration and cpudr operations
- 7034498 Incorrect return status on a failed CPUDR operation on a domain booted in single user mode.
- 7034787 Vdsdev device validation messages should be improved

- 7035102 Block 'ldm add-config' while 'ldm debug track' is in progress
- 7035296 Debug Logical Domains Manager hits an assertion failure when retrying a previously interrupted migration.
- 7035438 Primary switches to whole-core, but it does not have all vcpus of a core bound
- 7036440 HV mblock is modified at both ends
- 7036457 ldm rm-mem --auto-adj 256M primary failed on SPARC T3-1B
- 7036975 ldm migrate triggers free(1c5978): invalid or corrupted buffer
- 7037295 ldm add-spconfig -r <config> does not persist <config> over ldmd restarts
- 7037495 Sun Management Console operation dumps core in debug mode
- 7039020 Memory leaks found by parfait and ::findleaks
- 7039659 ldom_num_faulted_vcpus() hurts performance
- 7039793 With incoming_migration_enabled=false, outgoing migrations fail with version incompatibility message
- 7042294 LDoms MIB: memory leak in get_ldom_names
- 7042966 ldomCryptoResourcePool: ldomCryptoRpReserved property value is mismatch
- 7043129 ldmd needs an option to disable device validation
- 7044329 ldm add-vdisk/vnet/vcc/vsw/vds/vdpcs can result in devalias longer than OBP supports
- 7044716 ds_snmp: WARNING: Received DS snmp data out of sequence with request seen on T4
- 7046372 ldoms-dlmd: delayed reconfiguration produces a warning in the smf log
- 7046572 ldmd should not include 'power-management' node for guest md
- 7048883 Sun4v PRI needs a property to indicate max number of hv_mblock nodes
- 7049194 CLI: executing 'ldm add-domain with options' but no ldm name returns no syntax error
- 7049722 After installing redistributable svc:/ldoms/ldmd:default is in maintenance state on ldom
- 7050138 ldmd crashes on removing DRM policy from the primary in delayed reconfiguration mode.
- 7050588 install-ldm hangs when run via absolute path from other directories
- 7050996 Need the lowest numbered physical strand to stay in primary domain as workaround for 7043292

- 7053100 Python test failures on T4 systems with no MAUs
- 7053236 dr_vio.py fails occasionally in _testintervnetlink()
- 7053487 ldm man pages says 'stop' issues 'shutdown', which it does not
- 7054326 The hotplug command failed to disable/poweroff a PEM (PCI ExpressModule)
- 7056214 ldmd spews warnings "No delayed reconfig in progress" to ldmd log
- 7056745 mig_conn_t leaked in mig_listen_on_port()
- 7056809 ldmd doesn't need to listen on legacy migration port
- 7056909 ldmd should send REM_RESOURCE event to PM post cpu unbind
- 7059206 Error message for Mapin table memory allocation failure should be improved.
- 7060710 Auditing changes breaks Logical Domains Manager
- 7061064 RFE: ldm create should have an option to specify the domain UUID
- 7061114 Config changes induced by memory DR do not result in the current config switching to [next-poweron].
- 7063007 Double free in be_hv.c
- 7063470 RFE: Logical Domains Manager should dynamically figure out HV MD sizes and Guest MD sizes
- 7065441 ldm list sometimes prints 10.0% instead of 10%
- 7066070 basicSuite.py: "warning: Unknown soft state 0 for LDom ldg1" messages
- 7066447 WARNING: lgrp_minlat_node: malformed MD, no CPUs found in latency group
- 7066452 XMLv3 policy output doesn't show "disabled" in the elastic mode for the domain without whole-core
- 7066460 ldmp2v should handle missing s2 partition
- 7068869 CLI for add/set/rm-core to work differently from add/set/rm-vcpu -c
- 7069952 cancel-reconf restore the removed vdsdev
- 7070623 ldmd drops core after "set-vsw inter-vnet-link=off"
- 7070629 Domain migration checks for available console ports only on the first vcc on the target control domain
- 7071284 P2V: Add support allowing the transfer of a system which has zones
- 7071526 sparse-core strand allocations/deallocations are seen with bits that lock-down strand 0.
- 7072849 In PM elastic mode, memory leak found in cpu_del()

- 7073318 RFE:Memory allocation should support pinning specific memory ranges
- 7074339 Constant CPU DR of whole cores & threading changes causes ldmd to dump core
- 7075734 Use "add or remove" for whole-core error messages.
- 7076313 Implement SSC CLIs per Functional Spec
- 7080014 RFE: refactor domain suspend/resume to be less migration-centric
- 7080092 ldmd dumps core in xmpp_close_conn()
- 7081175 Add manual pkg(5) dependency on system/ldoms
- 7082794 ldm should not allow removal of PCIe bus after a PCIe device on that bus has been removed
- 7085572 Problems with list-io after removing a bus
- 7086033 RFE: XML support for SSC CLIs
- 7092958 Bind of I/O domain fails, unexpected devices in list-constraints
- 7093344 Remove ldmconfig(1M) from ldomsmanager pkg(5) manifest
- 7093862 Cannot bind ldom even there are console port available in second vccs
- 7096532 RFE: Relax whole-core switching restrictions on bound or active domains when no cap is specified
- 7097552 Allow only one instance of each ldm SVR4 package be installed at a time
- 7098704 ldm set-vsw does not update VLAN params in database
- 7098929 Number of vcpus in the DB is out of sync with Logical Domains Manager after max-cores is decreased
- 7100409 RFE: ldmd should send an XML event when a domain is suspended/resumed during a migration
- 7100431 add-vdsdev with null entry for mpgroup option still creates mpgroup
- 7100933 XMLv3 list-bindings/list-constraints : Assertion failed:
is_uuid_valid(ldp->props.uuid)
- 7102428 Primary mblock addresses laid out in some sequences cause endless system watchdog loops.
- 7102542 Domains with physical memory bindings encounter corruption during migration
- 7103183 Documentation should reflect that "set-vcpu -c" should not longer be used
- 7103270 ldmd loses UUID for control domain
- 7105688 Restrict ldmd -i option to debug mode

-
- 7105732 Add required dependencies to ldmd SMF service
 - 7106889 set-vsw doesn't update the MTU
 - 7107094 ldm rm-vdpcc doesn't update LDoms database
 - 7107548 Inactive domains can not be bound after applying patch 147507-01/-02
 - 7107827 Cancelling delayed reconfiguration does not restore extended-mapin-space value.
 - 7112190 XML vdisk section should include backend information
 - 7112205 XML should have a custom section
 - 7114415 Adding a very large number of mblocks to a domain causes ldmd SEGV crash and restart.
 - 7114920 ldm add-domain -i can fail but domain is created anyway
 - 7116203 Bump up the XML Interface version number from 1.2 to 1.3
 - 7116734 init-system prints the warning/failure message only from the first command processed
 - 7117114 Canceling delayed reconfigurations on systems not in del reconf does not return any messages.
 - 7117133 Primary panics on repeated delayed reconfig set/cancel of extended-mapin-space parameter.
 - 7117566 init-system does not restore named CPU cores when threading is max-ipc.
 - 7119504 Resource ID pools need better sparse ID allocation algorithm
 - 7120526 Disable use of 1/2 DR from PM code
 - 7120970 Tune the PM knobs for better responsiveness to load
 - 7123407 Last crypto unit is auto removed without -f option when last CPUs are DR'ed out via retries.
 - 7133885 ldmd dumps core (do_hv_pm_aync) on systems running S10 9/10
 - 7135380 Need to document impact of vdisk timeout on mpgroup feature
 - 7143201 Cross-CPU migration of guests with MAUs to platforms without MAUs fails
 - 7143721 Cannot create anymore a Virt. Switch with the java API, XMPP driver refuses the operation

