

Oracle® VM Server for SPARC 3.0 Release Notes

Copyright © 2007, 2013, 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 3.0 Release Notes	9
What's New in This Release	9
System Requirements	11
Supported Platforms	11
Required Software and Patches	12
Related Software	19
Optional Software	19
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 3.0 Software	21
Deprecated Features in the Oracle VM Server for SPARC 3.0 Release	21
Known Issues	22
General Issues	22
Domain Migration Restrictions	32
Oracle VM Server for SPARC MIB Issues	33
Bugs Affecting the Oracle VM Server for SPARC 3.0 Software	34
Documentation Errata	79
Resolved Issues	80

Preface

Oracle VM Server for SPARC 3.0 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 3.0 software.

Related Documentation

The following table shows the documentation that is available for the Oracle VM Server for SPARC 3.0 release. These documents are available in HTML and PDF formats unless indicated.

TABLE P-1 Related Documentation

Application	Title
Oracle VM Server for SPARC 3.0 Software	<i>Oracle VM Server for SPARC 3.0 Administration Guide</i> <i>Oracle VM Server for SPARC 3.0 Security Guide</i> <i>Oracle VM Server for SPARC 3.0 Reference Manual</i> <i>Oracle VM Server for SPARC 3.0 Release Notes</i>
Oracle VM Server for SPARC 3.0 drd(1M) and vntsd(1M) man pages	Oracle Solaris OS Reference Manuals: <ul style="list-style-type: none">■ Oracle Solaris 10 Documentation■ Oracle Solaris 11.1 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 11.1 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 field 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	Description	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 UNIX system prompts and superuser prompts for shells that are included in the Oracle Solaris OS. In command examples, the shell prompt indicates whether the command should be executed by a regular user or a user with privileges.

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 3.0 Release Notes

These release notes contain changes for this release, a list of supported platforms, a matrix of required software and patches, and other pertinent information, including bugs that affect the Oracle VM Server for SPARC 3.0 software.

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

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

What's New in This Release

The major changes for this release of the Oracle VM Server for SPARC 3.0 software are as follows. Note that each feature is available on all supported platforms unless otherwise stated.

- Adds support for Oracle's SPARC M5-32 platform. See “[Supported Platforms](#)” on page 11.
- Adds support for Oracle's SPARC T5 platform. See “[Supported Platforms](#)” on page 11.
- Adds support for the Fujitsu M10 system. See “[Supported Platforms](#)” on page 11.

- Enhances the resource affinity capability to address memory affinity.

This feature is only supported on Oracle's SPARC T3, SPARC T4, SPARC T5, and SPARC M5 platforms, and the Fujitsu M10 system.

- Retrieves a service processor (SP) configuration from the SP into the bootset, which is on the control domain. See [Chapter 11, “Managing Domain Configurations,”](#) in *Oracle VM Server for SPARC 3.0 Administration Guide*.

This feature is supported on the SPARC T3, SPARC T4, SPARC T5, and SPARC M5 platforms. Some limited support is also available on Oracle's UltraSPARC T2 and UltraSPARC T2 Plus platforms, and the Fujitsu M10 system.

- Adds cross-CPU migration support for Oracle Solaris 10. This feature enables you to perform a migration even if the CPU type of the source and target systems is not identical. Use the `ldm set -domain` command to set the `cpu-arch` property.

This feature is supported on the SPARC T3, SPARC T4 platforms, and Fujitsu M10 systems.

- Enhances the domain shutdown process. This feature enables you to specify whether to perform a full shutdown (default), a quick kernel stop, or force a shutdown. See the [ldm\(1M\)](#) man page.
- Enhances the domain shutdown process to perform an ordered shutdown of domains based on domain priorities. See *Fujitsu M10 Appendix* and the [ldm\(1M\)](#) man page.

This feature is supported only on the Fujitsu M10 system.

- Adds support for capacity on demand. See *Fujitsu M10 Systems Domain Configuration Guide*.

This feature is supported only on the Fujitsu M10 system.

- Adds support to automatically detect and degrade faulty memory and CPU resources. See “Viewing Faulty Resources” and “Autoreplacement of Faulty CPUs” in *Fujitsu M10 Appendix* and the [ldm\(1M\)](#) man page.

This feature is supported only on the Fujitsu M10 system.

- Adds support to collect hypervisor dump data on Fujitsu M10 systems. See “Hypervisor Dump Collection” in *Fujitsu M10 Appendix* and the [ldm\(1M\)](#) man page.
- Adds domain console logging support. See “Logging Domain Consoles” in *Fujitsu M10 Appendix*.

This feature is available on all supported platforms.

- Adds Oracle Solaris 11 support for Oracle VM Server for SPARC Management Information Base (MIB). See [Chapter 16, “Using the Oracle VM Server for SPARC Management Information Base Software,”](#) in *Oracle VM Server for SPARC 3.0 Administration Guide*.
- Enables a migration process to be initiated without specifying a password on the target system. See the [ldm\(1M\)](#) man page.
- Enables the live migration feature while the source machine, target machine, or both have the power management (PM) elastic policy in effect. See [Chapter 9, “Migrating Domains,”](#) in *Oracle VM Server for SPARC 3.0 Administration Guide*.
- Enables the dynamic resource management (DRM) feature while the host machine has the PM elastic policy in effect. See “Using Dynamic Resource Management” in *Oracle VM Server for SPARC 3.0 Administration Guide*.

This feature is supported only on the SPARC T3, SPARC T4, SPARC T5, and SPARC M5 platforms, and the Fujitsu M10 system.

- 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>).

The Oracle VM Server for SPARC 3.0 software is supported on the following platforms:

- **Fujitsu M-Series servers**
 - Fujitsu M10 system (refer to *Fujitsu M10 System Product Notes*)
- **SPARC M-Series servers**
 - SPARC M5-32 server (refer to *SPARC M5-32 Server Product Notes*)
- **SPARC T5 servers**
 - SPARC T5-1B server (refer to *SPARC T5-1B Server Product Notes*)
 - SPARC T5-2 server (refer to *SPARC T5-2 Server Product Notes*)
 - SPARC T5-4 server (refer to *SPARC T5-4 Server Product Notes*)
 - SPARC T5-8 server (refer to *SPARC T5-8 Server Product Notes*)
- **SPARC T4 servers**
 - SPARC T4-1 server (refer to *SPARC T4-1 Server Product Notes*)
 - SPARC T4-2 server (refer to *SPARC T4-2 Server Product Notes*)
 - SPARC T4-4 server (refer to *SPARC T4-4 Server Product Notes*)
 - SPARC T4-1B server (refer to *SPARC T4-1B Server Product Notes*)
 - Netra SPARC T4-1 server (refer to *Netra SPARC T4-1 Server Product Notes*)
 - Netra SPARC T4-2 server (refer to *Netra SPARC T4-2 Server Product Notes*)

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

Required Software and Patches

This section lists the required software and patches for use with the Oracle VM Server for SPARC 3.0 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 Versions

To use all features of the Oracle VM Server for SPARC 3.0 software, the operating system (OS) on all domains should be at least the Oracle Solaris 10 8/11 OS or the Oracle Solaris 11.1 Support Repository Update 1 (SRU 1). This OS can be either an initial installation of or an upgrade to the Oracle Solaris 10 8/11 OS or the Oracle Solaris 11.1 OS.

The minimum version of the Oracle Solaris 11 OS for each domain type is Oracle Solaris 11.1. 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 3.0 functionality.

For information about the minimum Oracle Solaris OS versions for supported server platforms, see the data sheet for your server platform at <http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html>.

TABLE 1-1 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	
148322-07 (Oracle Solaris 10 1/13 ixgbe)	X	X	X	
148324-06 (Oracle Solaris 10 1/13 ixgbev)	X	X	X	
148888-01 (Oracle Solaris 10 1/13 kernel update)	X	X	X	
148889-01 (Oracle Solaris 10 1/13 kernel update)	X	X	X	
149644-01 (Oracle Solaris 10 1/13 fmd)	X	X	X	
150107-01 (Oracle Solaris 10 1/13 ds)	X	X	X	X

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

Required Software to Enable Oracle VM Server for SPARC 3.0 Features

To enable all the Oracle VM Server for SPARC 3.0 features, you must run the required system firmware versions on the following platforms:

UltraSPARC T2 server	Run at least version 7.4.2 of the system firmware.
UltraSPARC T2 Plus server	Run at least version 7.4.2 of the system firmware.
SPARC T3 server	Run at least version 8.2.1.b of the system firmware.
SPARC T4 server	Run at least version 8.2.1.b of the system firmware.

This firmware is preinstalled on the SPARC T4. The required firmware for the SPARC T5 platforms, SPARC M5 platforms, and Fujitsu M10 systems is preinstalled on your system. For information about the required Oracle Solaris OS version, see [“Required and Recommended Oracle Solaris OS Versions” on page 13](#).

Note – The system firmware for SPARC T3 and SPARC T4 systems has been updated to version 8.2.2 to address some migration issues.

Required and Recommended System Firmware Patches

To take advantage of all features of Oracle VM Server for SPARC 3.0, ensure that your SPARC T-Series server from Oracle runs at least these revisions of the following system firmware patches, which can be installed on Oracle Solaris 10 and Oracle Solaris 11.1 systems:

147307-07	Sun SPARC Enterprise T5120 and T5220 Servers
147308-05	Sun Blade T6320 Server Module
147309-06	Netra T5220 Server
147310-07	Sun SPARC Enterprise T5140 and T5240 Servers
147311-05	Sun SPARC Enterprise T5440 Server
147312-06	Sun Blade T6340 Server Module
147313-06	Netra T5440 Server
147314-06	Sun Netra T6340 Server Module
148816-02	SPARC T3-1 Server
148817-02	SPARC T3-2 Server
148818-02	SPARC T3-4 Server
148819-02	SPARC T3-1B Server
148820-03	Netra SPARC T3-1 Server
148821-02	Netra SPARC T3-1B Server
148822-03	SPARC T4-1 Server
148823-03	SPARC T4-2 Server

148824-03	SPARC T4-4 Server
148825-02	SPARC T4-1B Server
148826-03	Netra SPARC T4-1 Server
148827-02	Netra SPARC T4-2 Server
148828-02	Netra SPARC T4-1B Server

Minimum Version of Software Required

Running the Oracle VM Server for SPARC 3.0 software along with at least the minimum versions of the other software components is supported. In such a configuration, you might not get all the features of the Oracle VM Server for SPARC 3.0 software. For the best results in production environments, run with the recommended system firmware version described in “Required Software to Enable Oracle VM Server for SPARC 3.0 Features” on page 13, and with the Oracle Solaris 10 8/11 OS, the Oracle Solaris 11 OS, or the Oracle Solaris 11.1 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.

The Oracle VM Server for SPARC 3.0 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). For information about the minimum versions of the Oracle Solaris OS for supported server platforms, see the data sheet for your server platform at <http://www.oracle.com/technetwork/documentation/oracle-sparc-ent-servers-189996.html>.

- SPARC T4 servers: 8.1.1
- SPARC T3 servers: 8.1.x
- UltraSPARC T2 Plus servers: 7.4.x
- UltraSPARC T2 servers: 7.4.x

The required system firmware for the SPARC T5 platforms, SPARC M5 platforms, and Fujitsu M10 systems is preinstalled on your system. For the minimum Oracle Solaris OS version, see the data sheet for your SPARC T5, SPARC M5, and Fujitsu M10 systems.

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 <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

The PCIe SR-IOV feature is supported on the SPARC T3, SPARC T4, SPARC T5, and SPARC M5 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.

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

- **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)) or the Oracle Solaris 10 1/13 OS.
 - 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 `ixgbevf` virtual function driver patches, which are 148225 and 148227, respectively.

Location of the Oracle VM Server for SPARC 3.0 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-3_0.zip` package from My Oracle Support. See “[How to Download the Logical Domains Manager Software \(Oracle Solaris 10\)](#)” in *Oracle VM Server for SPARC 3.0 Administration Guide*.
- **Oracle Solaris 11 OS.** Obtain the `ldomsmanager` package from the Oracle Solaris 11 Support Repository or from My Oracle Support. See “[How to Upgrade to the Oracle VM Server for SPARC 3.0 Software \(Oracle Solaris 11\)](#)” in *Oracle VM Server for SPARC 3.0 Administration Guide*.

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

For Oracle Solaris 10, the `OVM_Server_SPARC-3_0.zip` file that you download contains the following:

- Oracle VM Server for SPARC 3.0 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 3.0 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-3_0/  
  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

Oracle VM Server for SPARC 3.0 Administration Guide, *Oracle VM Server for SPARC 3.0 Reference Manual*, and this document 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 3.0 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 3.0 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 16, “Using the Oracle VM Server for SPARC Management Information Base Software,”](#) in *Oracle VM Server for SPARC 3.0 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.

- **Oracle VM Manager** is a web-based user interface that you can use to manage the Oracle VM environment. Earlier versions of this user interface managed only the Oracle VM Server x86 software, but starting with Oracle VM Manager 3.2.1 and Oracle VM Server for SPARC 3.0, you can also manage the Oracle VM Server for SPARC software. For more information about Oracle VM Manager, see the [Oracle VM Documentation \(http://www.oracle.com/technetwork/documentation/vm-096300.html\)](http://www.oracle.com/technetwork/documentation/vm-096300.html).
- **SunVTS** functionality is available in the control domain and guest domains on certain 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 *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 *Oracle Explorer User's Guide* (<http://docs.oracle.com/cd/E19957-01/819-6613/819-6613.pdf>).
- **Oracle Solaris Cluster** software can be used in a guest domain with some restrictions. See the Oracle Solaris Cluster documentation for more information about any restrictions and about the Oracle Solaris Cluster software in general. 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/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 3.0 software:

- **Sun Integrated Lights Out Manager (ILOM) 3.0** is the system management firmware that you can use to monitor, manage, and configure SPARC T-Series and SPARC M5 systems. ILOM is preinstalled on these platforms and can be used on supported servers with the Oracle VM Server for SPARC 3.0 software enabled. Refer to *Sun Integrated Lights Out Manager 3.0 User's Guide* for features and tasks that are common to Oracle's Sun rack-mounted servers or blade servers that support ILOM. Other user documents present ILOM features and tasks that are specific to the server platform that you are using. You can find the ILOM platform-specific information within the documentation set that accompanies your system.
- **Netra Data Plane Software Suite** is a complete board software package solution. The software provides an optimized rapid development and runtime environment on top of multithread partitioning firmware for Sun CMT platforms. The Logical Domains Manager contains some `ldm` subcommands (`add-vdpcs`, `rm-vdpcs`, `add-vdpcc`, and `rm-vdpcc`) for use with this software. For more information about this software, see the (<http://docs.oracle.com/cd/E19282-01/>) documentation.
- **Fujitsu M10 Systems eXtended System Control Facility (XSCF)** is the system management firmware that you can use to monitor, manage, and configure Fujitsu M10 systems. XSCF is preinstalled on these systems and can be used with the Oracle VM Server

for SPARC 3.0 software enabled. For more information about this software, refer to *Fujitsu M10 Systems System Operation and Administration Guide* and *Fujitsu M10 Systems XSCF Reference Manual*.

Upgrading to the Oracle VM Server for SPARC 3.0 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 3.0 software. See “[How to Upgrade to the Oracle VM Server for SPARC 3.0 Software \(Oracle Solaris 10\)](#)” in *Oracle VM Server for SPARC 3.0 Administration Guide* or “[How to Upgrade to the Oracle VM Server for SPARC 3.0 Software \(Oracle Solaris 11\)](#)” in *Oracle VM Server for SPARC 3.0 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.

Logical Domains Manager 3.0 is *not* supported on UltraSPARC T1 systems.

Deprecated Features in the Oracle VM Server for SPARC 3.0 Release

The XML v2 interface was deprecated in the Oracle VM Server for SPARC 1.2 release and is no longer supported in this Oracle VM Server for SPARC release.

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

- The network interface functionality of an Oracle VM Server for SPARC vsw driver was deprecated in Oracle Solaris 11.1. The Oracle VM Server for SPARC vsw driver continues to provide virtual network switching functionality for guest domains. See “[Oracle Solaris 11 Networking Overview](#)” in *Oracle VM Server for SPARC 3.0 Administration Guide*.

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

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

Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 3.0 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 3.0 Administration Guide*.

Service Processor and System Controller Are Interchangeable Terms

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

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

If a service domain is running a version of Oracle Solaris 10 OS prior to Oracle Solaris 10 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, Solaris Volume Manager might be unable to find its configuration or to access its metadevices.

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

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

- 1 **Boot the guest domain.**
- 2 **Disable the `dev_id` feature of Solaris Volume Manager by adding the following lines to the `/kernel/dr/md.conf` file:**

```
md_dev_id_destroy=1;
md_keep_repl_state=1;
```

- 3 **Reboot the guest domain.**
After the domain has booted, the Solaris Volume Manager configuration and metadevices should be available.
- 4 **Check the Solaris Volume Manager configuration and ensure that it is correct.**
- 5 **Re-enable the Solaris Volume Manager `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 devid 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. The LDC limits are:

UltraSPARC T2, SPARC T3-1, SPARC T3-1B, SPARC T4-1, and SPARC T4-1B servers

The LDC limit is 512.

UltraSPARC T2 Plus, the other SPARC T3 and SPARC T4 servers, SPARC T5, SPARC M5, and Fujitsu M10 systems

The LDC limit is 768.

This limitation becomes an issue only on the control domain because the control domain has at least part, if not all, of the I/O subsystem allocated to it. This limitation 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 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 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 smaller than the recommended size. For recommended and minimum memory requirements for the Oracle Solaris 10 OS, see [“System Requirements and Recommendations” in Oracle Solaris 10 8/11 Installation Guide: Planning for Installation and Upgrade](#). For recommended and minimum memory requirements for the Oracle Solaris 11 OS, see [Oracle Solaris 11 Release Notes](#) and [Oracle Solaris 11.1 Release Notes](#).

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

The memory dynamic reconfiguration (DR) feature enforces 256-Mbyte alignment on the address and size of the memory involved in a given operation. See [“Memory Alignment” in Oracle VM Server for SPARC 3.0 Administration Guide](#).

Booting a Large Number of Domains

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

- Up to 256 on Fujitsu M10 systems
- Up to 128 on SPARC M5 systems per physical domain
- Up to 128 on SPARC T5 systems

- Up to 128 on SPARC T4 servers
- Up to 128 on SPARC T3 servers
- Up to 128 on UltraSPARC T2 Plus servers
- Up to 64 on UltraSPARC T2 servers

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

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

Memory and swap space usage increases in a guest domain when the vsw services used by the domain provides services to many virtual networks in multiple domains. This increase is due to the peer-to-peer links between all the vnet connected to the vsw. The service domain benefits from having extra memory. The recommended minimum is 4 Gbytes when running more than 64 domains. Start domains in groups of 10 or fewer and wait for them to boot before starting the next batch. The same advice applies to installing operating systems on domains. You can reduce the number of links by disabling inter-vnet channels. See “[Inter-Vnet LDC Channels](#)” in *Oracle VM Server for SPARC 3.0 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.

▼ How to Power Off a System With Multiple Active Domains

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

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

▼ How to Power Cycle the System

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

Because no other domains are bound, the firmware automatically power cycles the system before rebooting it. When the system restarts, it boots into the 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 power cycle, unless the variable updates are either initiated from OpenBoot firmware on the control domain or followed by saving the configuration to the SC.

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

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

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

- At the OpenBoot prompt
- Using the Oracle Solaris OS `eeeprom(1M)` command
- Using the Logical Domains Manager CLI (`ldm`)
- 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 3.0 software, variable updates do not persist as expected in a few cases:

- All methods of updating a variable persist across reboots of that domain. However, they do not persist across a power cycle of the system, unless a subsequent logical domain configuration is saved to the SC. 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 power cycle 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: **15375997**, **15387338**, **15387606**, and **15415199**.

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

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

Delayed Reconfiguration

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

Cryptographic Units

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

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

ldmp2v convert Command: VxVM Warning Messages During Boot

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

```

Boot device: disk0:a File and args:
SunOS Release 5.10 Version Generic 139555-08 64-bit
Copyright 1983-2009 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Hostname: 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

```

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 shown in the following example.

```

# 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 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.

Note – A guest domain that runs a network-sensitive application might experience a short delay or interrupt while the migration is in progress.

Version Restrictions for Migration

If you attempt to perform a live migration of a domain that is booted with at least the 8.4 version of the firmware to a system that runs an older firmware version, the migration fails. The failure occurs because of a hypervisor API mismatch between the newer and older firmware versions. In this instance, you will see the following message:

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

```
Domain Migration of LDom ldg1 failed
```

Note that you can perform a live migration of a domain that is booted on a system with versions of the firmware that are older than the 8.4 version to a system that runs at least the 8.4 version of the firmware.

SPARC M5 and SPARC T5 Support for `cpu-arch=generic` Migration

Bug ID 15805135: In Oracle VM Server for SPARC 3.0 you *cannot* migrate a domain from a different platform type (such as SPARC T2, SPARC T2 Plus, SPARC T3, or SPARC T4) to either a SPARC T5 or SPARC M5 platform. This type of migration cannot occur even if you set `cpu-arch=generic`.

For SPARC T5 and SPARC M5 platforms, you can perform *only* migration operations between systems of the same platform type such as from a SPARC T5-2 to a SPARC T5-8 or from one SPARC M5-32 to another SPARC M5-32. This restriction only applies to SPARC T5 and SPARC M5 platforms.

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 3.0 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.

snmpable Command Does Not Work With the Version 2 or Version 3 Option

Bug ID 15376861: You receive empty SNMP tables if you query the Oracle VM Server for SPARC MIB 2.1 software using the `snmpable` command with the `-v2c` or `-v3` option. The `snmpable` 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 3.0 Administration Guide*.

Bugs Affecting the Oracle VM Server for SPARC 3.0 Software

This section summarizes the bugs that you might encounter when using this version of the software. The most recent bugs are described first. Workarounds and recovery procedures are specified, if available.

send_mondo_set: timeout Panic Occurs When Using the `ldm stop` Command on a Guest Domain After Stress

Bug ID 16486383: Assign a PCI device or bus directly to a guest domain where the domain does not have a core assigned from the `/SYS/DCU` where the PCI card physically resides. Because the hypervisor resets PCI devices on behalf of guest domains, during each guest domain reboot there is a chance that a domain with cores on the DCU connected to the PCI device will panic. More PCI devices assigned to non-DCU-local guests increases the possibility of panics.

Workaround: Perform one of the following workarounds:

- Ensure that when you assign PCI devices to a guest domain that the card is located physically in the same DCU as the cores.
- Manually assign cores to allow for physical card placement flexibility.
As an example, for a PCI device on I0U0 (`pci_0` through `pci_15`), choose a core between 0 and 127, and allocate it to the domain.

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

View the system cores by using the following command:

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

For a PCI device on I0U1 (`pci_16` through `pci_31`), choose a core between 128 and 255. For a PCI device on I0U2 (`pci_32` through `pci_47`), choose a core between 256 and 383. For a PCI device on I0U3 (`pci_48` through `pci_63`), choose a core between 384 and 511.

SPARC T5 and SPARC M5: While SR-IOV Devices Are in Use, Attempts to Unbind or Remove Resource Hangs And Cannot Be Stopped by Using Ctrl-C

Bug ID 16426940: On a SPARC T5 or a SPARC M5 system that has an SR-IOV configuration, you might experience a hang that cannot be stopped by Ctrl-C. The hang rarely occurs when you use the `ldm unbind` or `ldm rm-io` command.

Workaround: Reboot the instance of the Oracle Solaris OS that runs on the primary domain. Also, reboot any guest domain that uses I/O resources that are shared by the primary domain.

Dynamically Adding a Faulted Strand to a Domain Might Result in a Panic

Bug ID 16301304: In certain circumstances, a core that has been faulted is not immediately blacklisted by the Logical Domains Manager. As a result, it is possible to add strands from the faulted core to a domain. If such strands are added by using virtual CPU dynamic reconfiguration, the guest domain panics with a message similar to the following:

```
panic[cpu10]/thread=2a1003e9c60: promif_start_cpu: failed to start cpu 12 (6)
```

If a faulted strand is added to a domain that is not running, the following messages appears on the console when the domain is started:

```
NOTICE: cpux is not runnable and will not be brought online
NOTICE: cpux removed from system
```

These messages also appear on any subsequent reboot. While these messages do not cause any problems in the running domain, the listed CPUs will not be available to the guest domain.

Workaround: To avoid these problems, isolate the faulted strands from all domains. If the faulted strands are not in use, they cannot cause any adverse affects for any other domain in the system.

Subdevices Under a PCIe Device Revert to an Unassigned Name

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

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

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

SPARC M5-32: panic: mpo_cpu_add: Cannot read MD

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

```
mpo_cpu_add: Cannot read MD
```

The following procedure causes the panic:

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

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

The configuration successfully powers on and guest domains are booted.

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

Logical Domains generates a new GMD that reflects correct latency information, but the Oracle Solaris OS cannot parse the new information and panics.

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

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

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

SPARC M5-32: Root Complex Failover Might Lead to an Incorrect Direct I/O Guest Domain Configuration

Bug ID 16232834: The SPARC M5 platform includes dual internal ported disks and their dual paths. A root complex failover might lead to an incorrect direct I/O guest domain configuration.

Workaround: Ensure that both odd or both even PCIe Express Module (EMS) cards are assigned to the same domain. For example, if a domain has both EMS1 and EMS3, then both paths to the disk reside within that domain. The same is true if the domain has both EMS0 and EMS2.

Perform the following steps:

1. Disable the root complex failover feature by using the ILOM.

```
-> set /HOSTx ioreconfigure=false
```

x can be a value of 0 through 3.

2. Use the `ldm add-io` command to assign both odd or even EMS cards from the same I/O unit (IOU) to a guest domain.

In this example two domain configuration unit (DCU) host configuration, the following commands assign EMS cards to the `ldg1` guest domain.

First, the `pci_40` and `pci_44` root complexes are removed from the control domain.

```
# ldm rm-io pci_40 primary
# ldm rm-io pci_44 primary
```

Then, the `pci_40` and `pci_42` root complexes are added to the `ldg1` guest domain.

```
# ldm add-io pci_40 ldg1
# ldm add-io pci_44 ldg1
```

Similarly, you can assign both the `pci_48` and `pci_52` root complexes or all four buses to the guest domain.

After making these changes, use the `ldm ls -io` command to view the updated configuration.

```
# ldm ls-io
NAME                                TYPE  BUS      DOMAIN  STATUS
----                                -
pci_32                              BUS   pci_32   primary
pci_33                              BUS   pci_33   primary
pci_34                              BUS   pci_34   primary
pci_35                              BUS   pci_35   primary
pci_36                              BUS   pci_36   primary
pci_37                              BUS   pci_37   primary
pci_38                              BUS   pci_38   primary
pci_39                              BUS   pci_39   primary
pci_40                              BUS   pci_40   primary
```

pci_41	BUS	pci_41	primary	
pci_42	BUS	pci_42	primary	
pci_43	BUS	pci_43	primary	
pci_44	BUS	pci_44	primary	
pci_45	BUS	pci_45	primary	
pci_46	BUS	pci_46	primary	
pci_47	BUS	pci_47	primary	
pci_48	BUS	pci_48	primary	
pci_49	BUS	pci_49	primary	
pci_50	BUS	pci_50	primary	
pci_51	BUS	pci_51	primary	
pci_52	BUS	pci_52	primary	
pci_53	BUS	pci_53	primary	
pci_54	BUS	pci_54	primary	
pci_55	BUS	pci_55	primary	
pci_56	BUS	pci_56	primary	
pci_57	BUS	pci_57	primary	
pci_58	BUS	pci_58	primary	
pci_59	BUS	pci_59	primary	
/SYS/IOU2/PCIE3	PCIE	pci_32	primary	OCC
/SYS/IOU2/EMS1/CARD/NET0	PCIE	pci_32	primary	OCC
/SYS/IOU2/EMS1/CARD/SCSI	PCIE	pci_32	primary	OCC
/SYS/IOU2/PCIE2	PCIE	pci_33	primary	OCC
/SYS/IOU2/PCIE5	PCIE	pci_34	primary	EMP
/SYS/IOU2/PCIE8	PCIE	pci_35	primary	EMP
/SYS/IOU2/PCIE11	PCIE	pci_36	primary	EMP
/SYS/IOU2/EMS3/CARD/NET0	PCIE	pci_36	primary	OCC
/SYS/IOU2/EMS3/CARD/SCSI	PCIE	pci_36	primary	OCC
/SYS/IOU2/PCIE10	PCIE	pci_37	primary	OCC
/SYS/IOU2/PCIE13	PCIE	pci_38	primary	OCC
/SYS/IOU2/PCIE16	PCIE	pci_39	primary	OCC
/SYS/IOU2/PCIE6	PCIE	pci_40	primary	EMP
/SYS/IOU2/EMS2/CARD/NET0	PCIE	pci_40	primary	OCC
/SYS/IOU2/EMS2/CARD/SCSI	PCIE	pci_40	primary	OCC
/SYS/IOU2/PCIE7	PCIE	pci_41	primary	EMP
/SYS/IOU2/PCIE4	PCIE	pci_42	primary	EMP
/SYS/IOU2/PCIE1	PCIE	pci_43	primary	OCC
/SYS/IOU2/PCIE14	PCIE	pci_44	primary	OCC
/SYS/IOU2/EMS4/CARD/NET0	PCIE	pci_44	primary	OCC
/SYS/IOU2/EMS4/CARD/SCSI	PCIE	pci_44	primary	OCC
/SYS/IOU2/PCIE15	PCIE	pci_45	primary	OCC
/SYS/IOU2/PCIE12	PCIE	pci_46	primary	EMP
/SYS/IOU2/PCIE9	PCIE	pci_47	primary	EMP
/SYS/IOU3/PCIE3	PCIE	pci_48	primary	EMP
/SYS/IOU3/EMS1/CARD/NET0	PCIE	pci_48	primary	OCC
/SYS/IOU3/EMS1/CARD/SCSI	PCIE	pci_48	primary	OCC
/SYS/IOU3/PCIE2	PCIE	pci_49	primary	OCC
/SYS/IOU3/PCIE5	PCIE	pci_50	primary	OCC
/SYS/IOU3/PCIE8	PCIE	pci_51	primary	EMP
/SYS/IOU3/PCIE11	PCIE	pci_52	primary	EMP
/SYS/IOU3/PCIE12	PCIE	pci_52	primary	EMP
/SYS/IOU3/EMS3/CARD/NET0	PCIE	pci_52	primary	OCC
/SYS/IOU3/EMS3/CARD/SCSI	PCIE	pci_52	primary	OCC
/SYS/IOU3/PCIE9	PCIE	pci_53	primary	OCC
/SYS/IOU3/PCIE10	PCIE	pci_53	primary	OCC
/SYS/IOU3/PCIE13	PCIE	pci_54	primary	EMP
/SYS/IOU3/PCIE14	PCIE	pci_54	primary	EMP
/SYS/IOU3/EMS4/CARD/NET0	PCIE	pci_54	primary	OCC
/SYS/IOU3/EMS4/CARD/SCSI	PCIE	pci_54	primary	OCC

/SYS/IOU3/PCIE15	PCIE	pci_55	primary	EMP
/SYS/IOU3/PCIE16	PCIE	pci_55	primary	EMP
/SYS/IOU3/PCIE6	PCIE	pci_56	primary	OCC
/SYS/IOU3/EMS2/CARD/NET0	PCIE	pci_56	primary	OCC

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

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

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

This panic occurs if the following conditions exist:

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

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

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

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

1. Generate a new SP configuration after the DCUs have been added.
For example, the following command creates a configuration called `new-config-more-dcus`:

```
# ldm add-config new-config-more-dcus
```
2. Shutdown the domain.
3. Stop the host.

```
-> stop /HOST
```
4. Start the host.

```
-> start /HOST
```

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

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

Workaround: None.

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

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

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

ILOM Power Capper Should Synchronize Adjustment Updates with /SYS/VPS Updates

Bug ID 16205895: In each SPARC M5 physical domain, power capping monitors that domain's power consumption sensor, `/SYS/VPS`, to determine when to adjust CPU frequencies to meet a power limit. This sensor is only updated every 20-30 seconds. However, the sensor is polled and power adjustments are made based on its value every few seconds. The stale power consumption values can cause spikes in the throttling and unthrottling the CPU frequencies as the system adjusts unnecessarily.

Workaround: Do not use power capping.

A Guest Domain Cannot Boot When IOV-Loaned Devices Are Disabled

Bug ID 16098592: When a PCIe physical function or virtual function device is loaned to a guest domain, you can use the ILOM CLI or fault management framework to mark the device as disabled in the dynamic database (DDB). When the system is next started (power-on, reset), the host configuration reflects the DDB status to the hypervisor by means of the machine descriptor (when PRI/Guest MD etc are created). The hypervisor disables those devices but the Oracle Solaris OS is unaware that the device is disabled.

When guest domain that contains that device is started, the domain hangs. The domain hangs because the Oracle Solaris instance that owns the PCIe fabric cannot bring up the device as the hypervisor has disabled the device. As a result, when OBP attempts to probe PCIe devices, the hypervisor blocks access which leaves the OBP to continuously retry.

Workaround: Remove the devices that transitioned into the UNK (unknown) state from that guest domain by performing the following steps:

1. Stop the guest domain.
2. Remove the devices that are in the unknown state from the guest domain.

Use the `ldm rm-io` command.

- Restart the guest domain.

mutex_enter: bad mutex Panic in primary Domain During a Reboot or Shutdown

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

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

```
000002a1075c3630 ldc:ldc_mem_rdwr_cookie+20 (c4005fa01c80,
c4004e2c2000,2a1075c37c8, 6c80000, 1, 0)
%l0-3: 00000000001356a4 0000000000136800 0000000000000380
00000000000002ff
%l4-7: 00000000001ad3f8 0000000000000004 00000000ffbf89c
0000c4005fa01c88
000002a1075c3710 vldc:i_vldc_ioctl_write_cookie+a4 (c4004c400030,
380,ffbf898, 100003, 0, 70233400)
%l0-3: 00000000006c80000 0000000000156dc8 0000000000000380
0000000000100003
%l4-7: 00000000702337b0 000002a1075c37c8 0000000000040000
0000000000000000
000002a1075c37f0 vldc:vldc_ioctl+1a4 (3101, c4004c400030,
ffbf898,c4004c400000, c4004c438030, 0)
%l0-3: 0000000000100003 0000000000000000 000000007b340400
0000c4004c438030
%l4-7: 0000c4004c400030 0000000000000000 0000000000000000
0000000000000000
000002a1075c38a0 genunix:fop_ioctl+d0 (c4004d327800, 0, fbf898,
100003,c4004384f718, 2a1075c3acc)
%l0-3: 00000000000003103 0000000000100003 000000000133ce94
0000c4002352a480
%l4-7: 0000000000000000 0000000000000002 00000000000000c0
0000000000000000
000002a1075c3970 genunix:ioctl+16c (3, 3103, fbf898, 3, 134d50, 0)
%l0-3: 0000c40040e00a50 000000000000c6d3 0000000000000003
0000030000002000
%l4-7: 0000000000000003 0000000000000004 0000000000000000
0000000000000000
```

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

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

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

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

```
Dec 11 04:27:54 warning: Dropping pf
pci@00/pci@1/pci@0/pci@0/pci@4/LSI,sas@0: no IOV capable driver
```

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

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

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

Cannot Set a Jumbo MTU for sxge Virtual Functions In the primary Domain of a SPARCT5-1B

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

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

Previously Blacklisted Resources Are Erroneously Marked as Available With a Pending Evacuation After a Logical Domains Manager Restart

Bug ID 16016576: If the Logical Domains Manager is restarted, either manually or if the control domain is rebooted, any faulted resources that have been blacklisted become available. These resources are erroneously marked as available with a pending evacuation even though they had been previously evacuated and successfully added to the black list. The Logical Domains Manager does not prevent such resources from being added to a domain.

Recovery: No recovery is required. The evacuated resources are not assigned to any domain and therefore do not affect the current system configuration. However, because these resources were previously faulted, take care to avoid assigning these resources to any domain in the system.

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

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

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

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

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

Additional Restrictions for Migration

Bug ID 15858731: For Fujitsu M10 systems, the following restriction supersedes the information that is described in “[Migrating a Domain From the OpenBoot PROM or a Domain That Is Running in the Kernel Debugger](#)” in *Oracle VM Server for SPARC 3.0 Administration Guide*.

When a domain to be migrated is running in OpenBoot or in the kernel debugger (kmdb), the migration attempt always fails if the source machine or the target machine is a Fujitsu M10 system. In the case where the domain to be migrated has only a single CPU, you might see the following error message:

```
# ldm migrate ldg1 system2
Non-cooperative migration is not supported on this platform.
```

ldmconfig Fails: Fix for Bug ID 15972394 Precludes Unmodified factory-default Configuration

Bug ID 15829698: The `ldmconfig` command does not work with the Oracle VM Server for SPARC 3.0 software because it runs only when the system runs an unmodified factory-default configuration. A bug introduced in Oracle VM Server for SPARC 3.0 always reports the configuration as being modified.

Workaround: Instead of performing an initial installation of Oracle VM Server for SPARC 3.0, first install Oracle VM Server for SPARC 2.2 and then run `ldmconfig`. After the domains are created, upgrade the `SUNWldm` package to Oracle VM Server for SPARC 3.0.

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

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

For example:

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

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

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

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

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

```
# ldm ls
```

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

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

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

```
# ldm ls
```

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

Note that 760 exceeds the recommended maximum.

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

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

Recovery: Restart the logical domain.

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

SPARC T4-4: Unable to Bind a Guest Domain

Bug ID 15825330: Oracle VM Server for SPARC 3.0 appears to hang at startup on some SPARC T4-4 configurations that have only a single processor board.

Workaround: Ensure that a processor board always occupies the slots for processors 0 and 1. Restarting the system in such a configuration enables the Oracle VM Server for SPARC 3.0 software to start up.

Guest Domain Can Panic When CPU DR is Performed After a Cross-CPU Migration

Bug ID 15825060: When you use live migration to migrate a domain that was created on a SPARC T3 or SPARC T4 system to an UltraSPARC T2 or UltraSPARC T2 Plus system, a subsequent CPU dynamic reconfiguration operation might result in a panic. The panic message is similar to the following:

```
panic[cpu8]/thread=2a102491c60: cpu8: dev_mondo queue
configuration failed, error 6
```

Recovery: None.

Workaround: When migrating a domain that was created on a SPARC T3 or SPARC T4 system to an UltraSPARC T2 or UltraSPARC T2 Plus system, do not migrate an active domain. Instead, stop the domain before you begin the migration process.

Guest Domain Fails to Boot When the Memory is Highly Fragmented

Bug ID 15824270: A domain that runs the Oracle Solaris 11.1 OS and that has highly fragmented memory assignments might be unable to boot, and the following error is shown:

```
ERROR: Last Trap: Fast Data Access MMU Miss
```

Workaround: First, attempt to vary the amount of memory that is bound to the guest that fails to boot. If that attempt fails or if the primary domain is affected, perform a power cycle into an alternate SP configuration.

CPU Dynamic Reconfiguration of Available CPUs Causes an lpl_topo_verify fail -5 Panic

Bug ID 15823255: A domain on a SPARC M5 that has two more more DCUs might experience a panic under the following conditions:

- The domain must have at least two DCUs configured
- Only one DCU has CPUs configured
- The other DCUs must have no CPUs configured but must have memory configured
- Any one of the following operations cause the panic:
 - Turn CPUs on or off
 - Perform a dynamic reconfiguration (DR) of CPUs
 - Create or remove a processor set
 - Add or remove processors from a processor set

You can see the panic message on the console and in the `/var/adm/messages` file when the system reboots:

```
panlc [cpu4]/thread=0x30012a008: : lpl_topo_verify failed: -5
```

Workaround: Perform the following steps:

1. Refrain from performing any of the operations that might cause the panic.
2. Append the following line to the `/etc/system` file:

```
set lgrp_topo_levels=2
```

3. Reboot the system.

After the system reboots, you can safely perform the operations that previously caused the panic.

Root Domains Cannot Have Dependencies on Other Root Domains

Bug ID 15823203: A root domain that owns a PCIe bus *cannot* have its PCIe endpoint devices or SR-IOV virtual functions assigned to another root domain. However, you *can* assign a PCIe endpoint device or virtual function from a PCIe bus to the root domain that owns that bus.

Guest Domain Panics After a Migration When Cores are Added

Bug ID 15822313: On a system that runs the Oracle Solaris 11.1 OS, executing CPU DR operations on a domain that has been migrated can lead to a guest domain panic.

Workaround: Do not execute domain manager CPU DR operations until the migrated guest domain is rebooted.

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

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

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

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

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

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

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

ldm rm -io Should Accept --dry-run as an Alias for -n

Bug ID 15818302: You cannot specify the `--dry-run` option for the `ldm rm -io` command.

Workaround: Use the `-n` option instead.

SPARC T3 and SPARC T4: Domain With a Virtual Network Device or a Virtual Switch Under Heavy Stress Might Panic

Bug ID 15816287: In rare circumstances, a logical domain might panic if its virtual network device or virtual switch is under heavy stress.

Workaround: Perform one of the workarounds. It is best to implement the `/etc/system` workaround to limit the system performance impact.

- Add the following line to the `/etc/system` file for any domain that has a virtual switch and then reboot the domain for the changes to take affect:

```
set vsw:vsw_ldc_tx_serialize=1
```

- Set the `extended-mapin-space` property value to `off` for any domain that has a virtual network device or a switch.

```
# ldm set-domain extended-mapin-space=off ldom
```

Migration Problems Arise Between Systems That Have Different Firmware Versions Installed

Bug ID 15815409: The migration of an active domain might fail if the source machine and target machine run different system firmware versions. Such failure occurs in the following situations:

- **SPARC T3 or SPARC T4.** The source machine runs a firmware release older than 8.1.5 and the target machine runs at least 8.1.5.
- **UltraSPARC T2 or UltraSPARC T2 Plus.** The source machine runs a firmware release older than 7.4.3 and the target machine runs at least 7.4.3.

Recovery: Do one of the following:

- Install the appropriate system firmware on the system that runs the older version of the system firmware:
 - **SPARC T3 or SPARC T4.** Install at least system firmware version 8.1.5.
 - **UltraSPARC T2 or UltraSPARC T2 Plus.** Install at least system firmware version 7.4.3.
- If you cannot upgrade the system that runs the older firmware, install system firmware version 8.2.2.b on the system that runs the newer firmware.

Workaround: Stop the domain before you perform the migration.

No Error Message When a Memory DR Add is Partially Successful

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

Workaround: None.

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

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

Recovery: Restart the affected domain.

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

Unable to Stop a Domain With the `ldm stop` or `ldm stop -f` Command

Bug ID 15801579: In rare circumstances, it is not possible to stop a domain with either the `ldm stop` or the `ldm stop -f` command. This problem usually occurs only when some other issue causes the domain to either hard hang or enter a very rapid panic-reboot loop.

Workaround: If you cannot log in to the domain directly, attempt to issue the `ldm stop -f` command several times. This command might be executed in a time window during which the hypervisor can stop the domain.

If this workaround does not succeed, perform a power cycle of the system.

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

Bug ID 15789903: 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
```


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

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

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

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

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

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

Workaround: 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 49](#).

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

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

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

Workaround: You can ignore this message.

Ldm init-system Command Cannot Correctly Re-create a Domain That Has Virtual Function Devices

Bug ID 15783031: 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 15782994: Logical Domains Manager might crash and restart when you attempt an operation that affects the configuration of many domains. You might see this issue when you attempt to change anything related to the virtual networking configuration and if many virtual network devices in the same virtual switch exist across many domains. Typically, this issue is seen with approximately 90 or more domains that have virtual network devices connected to the same virtual switch, and the `inter-vnet-link` property is enabled (the default behavior). Confirm the symptom by finding the following message in the `ldmd` log file and a core file in the `/var/opt/SUNW/ldm` 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.

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

Bug ID 15780217: 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 15778392: The control domain requires the lowest core in the system. So, if core ID 0 is the lowest core, it *cannot* be shared with any other domain if you want to apply the whole-core constraint to the control domain.

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

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

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

ldmd Daemon Does Not Come Online

Bug ID 15777490: 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 older than the Oracle Solaris 10 10/09 OS. 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 15776752: If you cancel a live migration, the memory contents of the domain instance that is created on the target must be “scrubbed” by the hypervisor. This scrubbing process is performed for security reasons and must be complete before the memory can be returned to the pool of free memory. While this scrubbing is in progress, ldm commands become unresponsive. As a result, the Logical Domains Manager appears to be hung.

Recovery: You must wait for this scrubbing request to 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 15776319: On a system that runs the Oracle Solaris OS on the control domain and an I/O domain, some Emulex cards that are assigned to the I/O domain do not function properly because the cards do not receive interrupts. However, when assigned to the control domain, the same cards work properly.

This problem occurs with the following Emulex cards:

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

Workaround: None.

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

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

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

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

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

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

Workaround: You can ignore these misleading messages.

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

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

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

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

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

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

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

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

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

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

Workaround: Reboot the system.

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

Bug ID 15772090: 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
```

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

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

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

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

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

Bug ID 15761509: 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 15759601: If you issue an `ldm stop` command immediately after an `ldm start` command, the `ldm stop` command might fail with the following error:

```
LDom domain stop notification failed
```

Workaround: Reissue the `ldm stop` command.

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

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

Workaround: Perform the following steps:

1. Create an XML file for the primary domain.

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

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

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

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

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

**Bug ID 15753523:** 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 not to be set correctly 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 15750727:** 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 15748555:** 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 or Fujitsu M10 System Might Panic at Boot Time

**Bug ID 15748357:** 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 or Fujitsu M10 system. This problem 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 3.0 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 15748348:** When the primary domain shares the lowest physical core (usually 0) with another domain, attempts to set the whole-core constraint for the primary domain fail.

**Workaround:** Perform the following steps:

1. Determine the lowest bound core that is shared by the domains.

```
ldm list -o cpu
```

2. Unbind all the CPU threads of the lowest core from all domains other than the primary domain.

As a result, CPU threads of the lowest core are *not* shared and are free for binding to the primary domain.

3. Set the whole-core constraint by doing one of the following:

- Bind the CPU threads to the primary domain, and set the whole-core constraint by using the `ldm set -vcpu -c` command.
- Use the `ldm set -core` command to bind the CPU threads and set the whole-core constraint in a single step.

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

**Bug ID 15747047:** 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 has only 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.

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

**Bug ID 15738561:** The `ldm list -io` command might show the UNK or INV state for PCIe slots and SR-IOV virtual functions if the command runs immediately after the primary domain is booted. This problem is caused by the delay in the Logical Domains agent reply from the Oracle Solaris OS.

This problem has been only 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.

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

**Bug ID 15731303:** Avoid migrating domains that have over 500 Gbytes of memory. Use the `ldm list -o mem` command to see the memory configuration of your domain. Some memory configurations that have multiple memory blocks that total over 500 Gbytes might panic with a stack that resembles the following:

```
panic[cpu21]/thread=2a100a5dca0:
BAD TRAP: type=30 rp=2a100a5c930 addr=6f696e740a232000 mmu_fsr=10009

sched:data access exception: MMU sfsr=10009: Data or instruction address
out of range context 0x1

pid=0, pc=0x1076e2c, sp=0x2a100a5c1d1, tstate=0x4480001607, context=0x0
g1-g7: 80000001, 0, 80a5dca0, 0, 0, 0, 2a100a5dca0

000002a100a5c650 unix:die+9c (30, 2a100a5c930, 6f696e740a232000, 10009,
2a100a5c710, 10000)
000002a100a5c730 unix:trap+75c (2a100a5c930, 0, 0, 10009, 30027b44000,
2a100a5dca0)
000002a100a5c880 unix:ktl0+64 (7022d6dba40, 0, 1, 2, 2, 18a8800)
000002a100a5c9d0 unix:page_trylock+38 (6f696e740a232020, 1, 6f69639927eda164,
7022d6dba40, 13, 1913800)
000002a100a5ca80 unix:page_trylock_cons+c (6f696e740a232020, 1, 1, 5,
7000e697c00, 6f696e740a232020)
000002a100a5cb30 unix:page_get_mnode_freelist+19c (701ee696d00, 12, 1, 0, 19, 3)
000002a100a5cc80 unix:page_get_cachelist+318 (12, 1849fe0, ffffffff, 3,
0, 1)
000002a100a5cd70 unix:page_create_va+284 (192aec0, 300ddbc6000, 0, 0,
2a100a5cf00, 300ddbc6000)
000002a100a5ce50 unix:segkmem_page_create+84 (18a8400, 2000, 1, 198e0d0, 1000,
11)
000002a100a5cf60 unix:segkmem_xalloc+b0 (30000002d98, 0, 2000, 300ddbc6000, 0,
107e290)
000002a100a5d020 unix:segkmem_alloc_vn+c0 (30000002d98, 2000, 107e000, 198e0d0,
30000000000, 18a8800)
000002a100a5d0e0 genunix:vmem_xalloc+5c8 (30000004000, 2000, 0, 0, 80000, 0)
000002a100a5d260 genunix:vmem_alloc+1d4 (30000004000, 2000, 1, 2000,
30000004020, 1)
000002a100a5d320 genunix:kmem_slab_create+44 (30000056008, 1, 300ddbc4000,
18a6840, 30000056200, 30000004000)
000002a100a5d3f0 genunix:kmem_slab_alloc+30 (30000056008, 1, ffffffff,
0, 300000560e0, 30000056148)
000002a100a5d4a0 genunix:kmem_cache_alloc+2dc (30000056008, 1, 0, b9,
ffffffffffff, 2006)
000002a100a5d550 genunix:kmem_cpucache_magazine_alloc+64 (3000245a740,
3000245a008, 7, 6028f283750, 3000245a1d8, 193a880)
000002a100a5d600 genunix:kmem_cache_free+180 (3000245a008, 6028f2901c0, 7, 7,
7, 3000245a740)
000002a100a5d6b0 ldc:vio_destroy_mblks+c0 (6028efe8988, 800, 0, 200, 19de0c0, 0)
000002a100a5d760 ldc:vio_destroy_multipools+30 (6028f1542b0, 2a100a5d8c8, 40,
0, 10, 30000282240)
000002a100a5d810 vnet:vgen_unmap_rx_dring+18 (6028f154040, 0, 6028f1a3cc0, a00,
200, 6028f1abc00)
000002a100a5d8d0 vnet:vgen_process_reset+254 (1, 6028f154048, 6028f154068,
```

6028f154060, 6028f154050, 6028f154058)  
 000002a100a5d9b0 genunix:taskq\_thread+3b8 (6028ed73908, 6028ed738a0, 18a6840,  
 6028ed738d2, e4f746ec17d8, 6028ed738d4)

**Workaround:** Avoid performing migrations of domains that have over 500 Gbytes of memory.

## Removing a Large Number of CPUs From a Guest Domain Fails

**Bug ID 15726205:** You might see the following error message when you attempt to remove a large number of CPUs from a guest domain:

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

**Workaround:** Stop the guest domain before you remove more than 100 CPUs from the domain.

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

**Bug ID 15721872:** You *cannot* use Oracle Solaris hot-plug operations to hot-remove a PCIe endpoint device after that device is removed from the primary domain by using the `ldm rm-io` command. For information about replacing or removing a PCIe endpoint device, see [“Making PCIe Hardware Changes” in Oracle VM Server for SPARC 3.0 Administration Guide](#).

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

**Bug ID 15713809:** 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
```

Install 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 for both issues is 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 15710957:** When a heavily loaded guest domain has a hybrid I/O configuration and you attempt to migrate it, you might see an `nxge` panic.

**Workaround:** Add the following line to the `/etc/system` file on the primary domain and on any service domain that is part of the hybrid I/O configuration for the domain:

```
set vsw:vsw_hio_max_cleanup_retries = 0x200
```

## **All ldm Commands Hang When Migrations Have Missing Shared NFS Resources**

**Bug ID 15708982:** An initiated or ongoing migration, or any `ldm` command, hangs forever. This situation occurs when the domain to be migrated uses a shared file system from another system and the file system is no longer shared.

**Workaround:** Make the shared file system accessible again.

## Logical Domains Agent Service Does Not Come Online If the System Log Service Does Not Come Online

**Bug ID 15707426:** If the system log service, `svc:/system/system-log`, fails to start and does not come online, the Logical Domains agent service will not come online. When the Logical Domains agent service is not online, the `virtinfo`, `ldm add-vsw`, `ldm add-vdsdev`, and `ldm list-io` commands might not behave as expected.

**Workaround:** Ensure that the `svc:/ldoms/agents:default` service is enabled and online:

```
svcs -l svc:/ldoms/agents:default
```

If the `svc:/ldoms/agents:default` service is offline, verify that the service is enabled and that all dependent services are online.

## Kernel Deadlock Causes Machine to Hang During a Migration

**Bug ID 15704500:** The migration of an active guest domain might hang and cause the source machine to become unresponsive. When this problem occurs, the following message is written to the console and to the `/var/adm/messages` file:

```
vcc: i_vcc_ldc_fini: cannot close channel 15
vcc: [ID 815110 kern.notice] i_vcc_ldc_fini: cannot
close channel 15
```

Note that the channel number shown is an Oracle Solaris internal channel number that might be different for each warning message.

**Workaround:** Before you migrate the domain, disconnect from the guest domain's console.

**Recovery:** Perform a power cycle of the source machine.

## DRM and `ldm list` Output Shows a Different Number of Virtual CPUs Than Are Actually in the Guest Domain

**Bug ID 15702475:** A No response message might appear in the Oracle VM Server for SPARC log when a loaded domain's DRM policy expires after the CPU count has been substantially reduced. The `ldm list` output shows that more CPU resources are allocated to the domain than is shown in the `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 15701865:** If you attempt a live migration of a domain that depends on an inactive domain on the target machine, the `ldmd` daemon faults with a segmentation fault, and the domain on the target machine restarts. Although you can still perform a migration, it will not be a live migration.

**Workaround:** Perform one of the following actions before you attempt the live migration:

- Remove the guest dependency from the domain to be migrated.
- Start the master domain on the target machine.

## DRM Fails to Restore the Default Number of Virtual CPUs for a Migrated Domain When the Policy Is Removed or Expired

**Bug ID 15701853:** After you perform a domain migration while a DRM policy is in effect, if the DRM policy expires or is removed from the migrated domain, DRM fails to restore the original number of virtual CPUs to the domain.

**Workaround:** If a domain is migrated while a DRM policy is active and the DRM policy is subsequently expired or removed, reset the number of virtual CPUs. Use the `ldm set -vcpu` command to set the number of virtual CPUs to the original value on the domain.

## Virtual CPU Timeout Failures During DR

**Bug ID 15701258:** Running the `ldm set -vcpu 1` command on a guest domain that has over 100 virtual CPUs and some cryptographic units fails to remove the virtual CPUs. The virtual CPUs are not removed because of a DR timeout failure. The cryptographic units are successfully removed.

**Workaround:** Use the `ldm rm -vcpu` command to remove all but one of the virtual CPUs from the guest domain. Do *not* remove more than 100 virtual CPUs at a time.

## Migration Failure Reason Not Reported When the System MAC Address Clashes With Another MAC Address

**Bug ID 15699763:** A domain cannot be migrated if it contains a duplicate MAC address. Typically, when a migration fails for this reason, the failure message shows the duplicate MAC address. However in rare circumstances, this failure message might not report the duplicate MAC address.



```
ldm migrate ldg2 system2
Target Password:
Domain Migration of LDom ldg2 failed
```

**Workaround:** Ensure that the MAC addresses on the target machine are unique.

## Simultaneous Migration Operations in “Opposite Direction” Might Cause ldm to Hang

**Bug ID 15696986:** If two `ldm migrate` commands are issued simultaneously in the “opposite direction,” the two commands might hang and never complete. An opposite direction situation occurs when you simultaneously start a migration on machine A to machine B and a migration on machine B to machine A.

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

**Workaround:** None.

## Removing a Large Number of CPUs From the Control Domain Fails

**Bug ID 15677358:** Use a delayed reconfiguration rather than dynamic reconfiguration to remove more than 100 CPUs from the control domain (also known as the primary domain). Use the following steps:

1. Use the `ldm start -reconf primary` command to put the control domain in delayed reconfiguration mode.
2. Remove the desired number of CPU resources.

If you make a mistake while removing CPU resources, do *not* attempt another request to remove CPUs while the control domain is still in a delayed reconfiguration. If you do so, the commands will fail (see [“Only One CPU Configuration Operation Is Permitted to Be Performed During a Delayed Reconfiguration” on page 32](#)). Instead, undo the delayed reconfiguration operation by using the `ldm cancel -reconf` command, and start over.

3. Reboot the control domain.

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

**Bug IDs 15672651 and 15731467:** You might experience OS hangs at login or while executing commands when the following conditions are met:

- The Oracle Solaris 10 8/11 OS is running on a SPARC sun4v system
- The Power Management (PM) Elastic policy is set on the system's ILOM Service Processor

**Workaround:** Apply patch ID 147149-01.

## **pkgadd Fails to Set ACL Entries on /var/svc/manifest/platform/sun4v/ldmd.xml**

**Bug ID 15668881:** When using the `pkgadd` command to install the `SUNWldm.v` package from a directory that is exported by means of NFS from a Sun ZFS Storage Appliance, you might see the following error message:

```
cp: failed to set acl entries on /var/svc/manifest/platform/sun4v/ldmd.xml
```

**Workaround:** Ignore this message.

## **SPARCT3-1: Issue With Disks That Are Accessible Through Multiple Direct I/O Paths**

**Bug ID 15668368:** A SPARC T3-1 system can be installed with dual-ported disks, which can be accessed by two different direct I/O devices. In this case, assigning these two direct I/O devices to different domains can cause the disks to be used by both domains and affect each other based on the actual usage of those disks.

**Workaround:** Do not assign direct I/O devices that have access to the same set of disks to different I/O domains. To determine whether you have dual-ported disks on T3-1 system run the following command on the SP:

```
-> show /SYS/SASBP
```

If the output includes the following `fru_description` value, the corresponding system has dual-ported disks:

```
fru_description = BD,SAS2,16DSK,LOUISE
```

If dual disks are found to be present in the system, ensure that both of the following direct I/O devices are always assigned to the same domain:

```
pci@400/pci@1/pci@0/pci@4 /SYS/MB/SASHBA0
pci@400/pci@2/pci@0/pci@4 /SYS/MB/SASHBA1
```

## **Memory DR Removal Operations With Multiple Plumbed NIU nxge Instances Can Hang Indefinitely and Never Complete**

**Bug ID 15667770:** When multiple NIU `nxge` instances are plumbed on a domain, the `ldm rm-mem` and `ldm set-mem` commands, which are used to remove memory from the domain, might never complete. To determine whether the problem has occurred during a memory removal operation, monitor the progress of the operation with the `ldm list -o status` command. You might have encountered this problem if the progress percentage remains constant for several minutes.

**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 15664666:** When a reset dependency is created, an `ldm stop -a` command might result in a domain with a reset dependency being restarted instead of only stopped.

**Workaround:** First, issue the `ldm stop` command to the master domain. Then, issue the `ldm stop` command to the slave domain. If the initial stop of the slave domain results in a failure, issue the `ldm stop -f` command to the slave domain.

## Migration of a Domain That Has an Enabled Default DRM Policy Results in a Target Domain Being Assigned All Available CPUs

**Bug ID 15655513:** 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 15655199:** Sometimes an in-use MAC address is not detected and it is erroneously reassigned.

**Workaround:** Manually ensure that an in-use MAC address cannot be reassigned.

## Ldmconfig Cannot Create a Domain Configuration on the SP

**Bug ID 15654965:** The `ldmconfig` script cannot properly create a stored 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 15653424:** The migration of an active domain can fail if it is running a release earlier than the Oracle Solaris 10 10/09 OS release and the lowest-numbered CPU in the domain is in the `offLine` state. The operation fails when the Logical Domains Manager uses CPU DR to reduce the domain to a single CPU. In doing so, the Logical Domains Manager attempts to remove all but the lowest CPU in the domain but because that CPU is offline, the operation fails.

**Workaround:** Before attempting the migration, ensure that the lowest-numbered CPU in the domain is in the `onLine` state.

## Memory DR Is Disabled Following a Canceled Migration

**Bug ID 15646293:** After an Oracle Solaris 10 9/10 domain has been suspended as part of a migration operation, memory dynamic reconfiguration (DR) is disabled. This action occurs not only when the migration is successful but also when the migration has been canceled, despite the fact that the domain remains on the source machine.

## Dynamic Reconfiguration of MTU Values of Virtual Network Devices Sometimes Fails

**Bug ID 15631119:** If you modify the maximum transmission unit (MTU) of a virtual network device on the control domain, a delayed reconfiguration operation is triggered. If you subsequently cancel the delayed reconfiguration, the MTU value for the device is not restored to the original value.

**Recovery:** Rerun the `ldm set -vnet` command to set the MTU to the original value. Resetting the MTU value puts the control domain into delayed reconfiguration mode, which you need to cancel. The resulting MTU value is now the original, correct MTU value.

```
ldm set-vnet mtu=orig-value vnet1 primary
ldm cancel-op reconf primary
```

## Migrated Domain With MAUs Contains Only One CPU When Target OS Does Not Support DR of Cryptographic Units

**Bug ID 15606220:** Starting with the Logical Domains 1.3 release, a domain can be migrated even if it has one or more cryptographic units bound to it.

In the following circumstances, the target machine will contain only one CPU after the migration is completed:

- Target machine runs Logical Domains 1.2
- Control domain on the target machine runs a version of the Oracle Solaris OS that does not support cryptographic unit DR
- You migrate a domain that contains cryptographic units

After the migration completes, the target domain will resume successfully and be operational, but will be in a degraded state (just one CPU).

**Workaround:** Prior to the migration, remove the cryptographic unit or units from the source machine that runs Logical Domains 1.3.

**Mitigation:** To avoid this issue, perform one or both of these steps:

- Install the latest Oracle VM Server for SPARC software on the target machine.
- Install patch ID 142245-01 on the control domain of the target machine, or upgrade to at least the Oracle Solaris 10 10/09 OS.

## Confusing Migration Failure Message for Real Address Memory Bind Failures

**Bug ID 15605806:** In certain situations, a migration fails with the following error message, and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This situation can occur even if the total amount of available memory on the target machine is greater than the amount of memory being used by the source domain (as shown by `ldm ls-devices -a mem`).

```
Unable to bind 29952M memory region at real address 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 15600969:** If all the hardware cryptographic units are dynamically removed from a running domain, the cryptographic framework fails to seamlessly switch to the software cryptographic providers, and kills all the `ssh` connections.

**Recovery:** Re-establish the `ssh` connections after all the cryptographic units are removed from the domain.

**Workaround:** Set `UseOpenSSLEngine=no` in the `/etc/ssh/sshd_config` file on the server side, and run the `svcadm restart ssh` command.

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

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

**Bug ID 15597025:** When you run the `ldm ls-io -l` command on a system that has 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 15572184:** An `ldm` command might be slow to respond when several domains are booting. If you issue an `ldm` command at this stage, the command might appear to hang. Note that the `ldm` command will return after performing the expected task. After the command returns, the system should respond normally to `ldm` commands.

**Workaround:** Avoid booting many domains simultaneously. However, if you must boot several domains at once, refrain from issuing further `ldm` commands until the system returns to normal. For instance, wait for about two minutes on Sun SPARC Enterprise T5140 and T5240 servers and for about four minutes on the Sun SPARC Enterprise T5440 server or Netra T5440 server.

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

**Bug ID 15560811:** In Oracle Solaris 11, zones that are configured with an automatic network interface (`anet`) might fail to start in a domain that has Logical Domains virtual network devices only.

- **Workaround 1:** Assign one or more physical network devices to the guest domain. Use PCIe bus assignment, the Direct I/O (DIO), or the SR-IOV feature to assign a physical NIC to the domain.
- **Workaround 2:** If the zones configuration requirement is to have interzone communication only within the domain, create an `etherstub` device. Use the `etherstub` device as the “lower-link” in the zones configuration so that that virtual NICs are created on the `etherstub` device.
- **Workaround 3:** Use exclusive link assignment to assign a Logical Domains virtual network device to a zone. Assign virtual network devices, as needed, to the domain. You might also choose to disable inter-vnet links to be able to create a large number of virtual network devices.

## Oracle Solaris 10: Virtual Network Devices Are Not Created Properly on the Control Domain

**Bug ID 15560201:** Sometimes `ifconfig` indicates that the device does not exist after you add a virtual network or virtual disk device to a domain. This situation might occur as the result of the `/devices` entry not being created.

Although this problem should not occur during normal operation, the error 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 15555509:** When Logical Domains is configured on a system and you add another XAUI network card, the card is not visible after the machine has undergone a power cycle.

**Recovery:** To make the newly added XAUI visible in the control domain, perform the following steps:

1. Set and clear a dummy variable in the control domain.

The following commands use a dummy variable called `fix-xaui`:

```
ldm set-var fix-xaui=yes primary
ldm rm-var fix-xaui primary
```

2. Save the modified configuration to the SP, replacing the current configuration.

The following commands use a configuration name of `config1`:

```
ldm rm-spconfig config1
ldm add-spconfig config1
```

3. Perform a reconfiguration reboot of the control domain.

```
reboot -- -r
```

At this time, you can configure the newly available network or networks for use by Logical Domains.

## I/O Domain or Guest Domain Panics When Booting From e1000g

**Bug ID 15543982:** You can configure a maximum of two domains with dedicated PCI-E root complexes on systems such as the Sun Fire T5240. These systems have two UltraSPARC T2 Plus CPUs and two I/O root complexes.

`pci@500` and `pci@400` are the two root complexes in the system. The primary domain will always contain at least one root complex. A second domain can be configured with an unassigned or unbound root complex.

The `pci@400` fabric (or leaf) contains the on-board e1000g network card. The following circumstances could lead to a domain panic:

- If the system is configured with a primary domain that contains `pci@500` and a second domain that contains `pci@400`

---

**Note** – For some blades, the primary domain (system disk) is on the `pci@400` bus by default.

---

- The e1000g device on the `pci@400` fabric is used to boot the second domain

Avoid the following network devices if they are configured in a non-primary domain:

```
/pci@400/pci@0/pci@c/network@0,1
/pci@400/pci@0/pci@c/network@0
```

When these conditions are true, the domain will panic with a PCI-E Fatal error.

Avoid such a configuration or, if the configuration is used, do not boot from the listed devices.

## Explicit Console Group and Port Bindings Are Not Migrated

**Bug ID 15527921:** During a migration, any explicitly assigned console group and port are ignored, and a console with default properties is created for the target domain. This console is created using the target domain name as the console group and using any available port on the first virtual console concentrator (vcc) device in the control domain. If there is a conflict with the default group name, the migration fails.

**Recovery:** To restore the explicit console properties following a migration, unbind the target domain and manually set the desired properties using the `ldm set -vcons` command.

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

**Bug ID 15523133:** If the virtual disk on the target machine does not point to the same disk back end that is used on the source machine, the migrated domain cannot access the virtual disk using that disk back end. A hang can result when accessing the virtual disk on the domain.

Currently, the Logical Domains Manager checks only that the virtual disk volume names match on the source and target machines. In this scenario, no error message is displayed if the disk back ends do not match.

**Workaround:** 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:

- Migrate the domain and fix the `vdsdev`.
  1. Migrate the domain back to the source machine.
  2. Fix the `vdsdev` on the target to point to the correct disk back end.
  3. Migrate the domain to the target machine again.
- Stop and unbind the domain on the target, and fix the `vdsdev`. If the OS supports virtual I/O dynamic reconfiguration, and the incorrect virtual disk is not in use on the domain (that is, it is not the boot disk and is unmounted), do the following:
  1. Use the `ldm rm-vdisk` command to remove the disk.

2. Fix the vdsdev.
3. Use the `ldm add-vdisk` command to add the virtual disk again.

## Migration Can Fail to Bind Memory Even If the Target Has Enough Available

**Bug ID 15523120:** In certain situations, a migration fails and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This situation can occur even if the total amount of available memory on the target machine is greater than the amount of memory being used by the source domain.

This failure occurs because migrating the specific memory ranges in use by the source domain requires that compatible memory ranges are available on the target as well. When no such compatible memory range is found for any memory range in the source, the migration cannot proceed.

**Recovery:** If this condition is encountered, you might be able to migrate the domain if you modify the memory usage on the target machine. To do this, unbind any bound or active logical domain on the target.

Use the `ldm list-devices -a mem` command to see what memory is available and how it is used. You might also need to reduce the amount of memory that is assigned to another domain.

## Logical Domains Manager Does Not Start If the Machine Is Not Networked and an NIS Client Is Running

**Bug ID 15518409:** If you do not have a network configured on your machine and have a Network Information Services (NIS) client running, the Logical Domains Manager will not start on your system.

**Workaround:** Disable the NIS client on your non-networked machine:

```
svcadm disable nis/client
```

## Logical Domains Manager Displays Migrated Domains in Transition States When They Are Already Booted

**Bug ID 15516245:** On occasion, an active logical domain appears to be in the *transition* state instead of the *normal* state long after it is booted, or following the completion of a domain migration. This glitch is harmless, and the domain is fully operational. To see what flag is set, check the `flags` field in the `ldm list -l -p` command output, or check the `FLAGS` field in the `ldm list` command, which shows `-n----` for *normal* or `-t----` for *transition*.

**Recovery:** After the next reboot, the domain shows the correct state.

## Cannot Connect to Migrated Domain's Console Unless vntsd Is Restarted

**Bug ID 15513998:** Occasionally, after a domain has been migrated, it is not possible to connect to the console for that domain.

**Workaround:** Restart the vntsd SMF service to enable connections to the console:

```
svcadm restart vntsd
```

---

**Note** – This command will disconnect all active console connections.

---

## Sometimes, Executing the uadmin 1 0 Command From a Logical Domains System Does Not Return the System to the OK Prompt

**Bug ID 15511551:** Sometimes, executing the uadmin 1 0 command from the command line of a Logical Domains system does not leave the system at the ok prompt after the subsequent reset. This incorrect behavior is seen only when the Logical Domains variable auto-reboot? is set to true. If auto-reboot? is set to false, the expected behavior occurs.

**Workaround:** Use this command instead:

```
uadmin 2 0
```

Or, always run with auto-reboot? set to false.

## Logical Domains Manager Can Take Over 15 Minutes to Shut Down a Domain

**Bug ID 15505014:** A domain shutdown or memory scrub can take over 15 minutes with a single CPU and a very large memory configuration. During a shutdown, the CPUs in a domain are used to scrub all the memory owned by the domain. The time taken to complete the scrub can be quite long if a configuration is imbalanced, for example, a single CPU domain with 512 Gbytes of memory. This prolonged scrub time extends the amount of time needed to shut down a domain.

**Workaround:** Ensure that large memory configurations (>100 Gbytes) have at least one core.

## 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 15482406:** 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.

## scadm Command Can Hang Following an SC or SP Reset

**Bug ID 15469227:** 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 re-establish a connection following an SC reset.

**Workaround:** Reboot the host to re-establish connection with the SC.

**Recovery:** Reboot the host to re-establish connection with the SC.

## Simultaneous Netinstallation of Multiple Domains Fails When in a Common Console Group

**Bug ID 15453968:** Simultaneous net installation of multiple guest domains fails on systems that have a common console group.

**Workaround:** Only net-install on guest domains that each have their own console group. This failure is seen only on domains with a common console group shared among multiple net-installing domains.

## ldc\_close: (0xb) unregister failed, 11 Warning Messages

**Bug ID 15426914:** You might see the following warning message on the system console or in the system log:

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

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 15422900:** If you configure more than four virtual networks (vnets) in a guest domain on the same network using the Dynamic Host Configuration Protocol (DHCP), the guest domain can eventually become unresponsive while running network traffic.

**Workaround:** Set `ip_ire_min_bucket_cnt` and `ip_ire_max_bucket_cnt` to larger values, such as 32, if you have 8 interfaces.

**Recovery:** Issue an `ldm stop-domain ldom` command followed by an `ldm start-domain ldom` command on the guest domain (*ldom*) in question.

## OpenBoot PROM Variables Cannot be Modified by the `eeprom` Command When the Logical Domains Manager Is Running

**Bug ID 15387338:** This issue is summarized in “[Logical Domains Variable Persistence](#)” on [page 27](#) and affects only the control domain.

## Cannot Set Security Keys With Logical Domains Running

**Bug ID 15370442:** The Logical Domains environment, does not support setting or deleting wide-area network (WAN) boot keys from within the Oracle Solaris OS by using the `ickey(1M)` command. All `ickey` operations fail with the following error:

```
ickey: setkey: ioctl: I/O error
```

In addition, WAN boot keys that are set using OpenBoot firmware in logical domains other than the control domain are not remembered across reboots of the domain. In these domains, the keys set from the OpenBoot firmware are only valid for a single use.

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

**Bug ID 15368170:** In some cases, the behavior of the `ldm stop-domain` command is confusing.

```
ldm stop-domain -f ldom
```

If the domain is at the kernel module debugger, `kldb(1)`, prompt, then the `ldm stop-domain` command fails with the following error message:

```
LDom <domain-name> stop notification failed
```

## Documentation Errata

This section contains documentation errors that have been found too late to resolve for the Oracle VM Server for SPARC 3.0 release.

### **ldm(1M) Man Page: I/O Virtualization Commands Do Not Automatically Initiate a Delayed Reconfiguration**

Several places in the `ldm(1M)` man page state that the I/O virtualization commands automatically initiate a delayed reconfiguration. This is not true. If needed, you must manually initiate a delayed reconfiguration.

### **ldm(1M) Man Page: Dynamic Creation of Virtual Functions Is Not Supported**

The “Create a Virtual Function” section of the `ldm(1M)` man page states, “To dynamically create virtual functions, ensure that you set the `iov` property for the parent root complex.” This functionality is not supported in the Oracle VM Server for SPARC 3.0 release.

### **ldm(1M) Man Page: Only `ldm add-spconfig -r` Performs a Manual Recovery**

The description of the `-r` option in the `ldm(1M)` man page currently states that the `add-spconfig`, `list-spconfig`, and `remove-spconfig` subcommands use this option to perform a manual recovery. This is not true. Only the `ldm add-spconfig -r` command can be used to perform the manual recovery.

### **Additional Restrictions for Migration**

For Fujitsu M10 systems, the following restriction supersedes the information that is described in “[Migrating a Domain From the OpenBoot PROM or a Domain That Is Running in the Kernel Debugger](#)” in *Oracle VM Server for SPARC 3.0 Administration Guide*.

When a domain to be migrated is running in OpenBoot or in the kernel debugger (`kmdb`), the migration attempt always fails if the source machine or the target machine is a Fujitsu M10 system. In the case where the domain to be migrated has only a single CPU, you might see the following error message:

```
ldm migrate ldg1 system2
Non-cooperative migration is not supported on this platform.
```

## Resolved Issues

The following RFEs and bugs have been fixed for the Oracle VM Server for SPARC 3.0 software release:

- 6773569 RFE: Constraint database needs to be tied to saved config
- 6885505 RFE: Domain Migration should work in PM Elastic mode
- 885507 RFE: DRM should work in PM Elastic mode
- 908016 RFE: ldm stop should trigger a graceful Solaris shutdown
- 963878 RFE: Need a way to retrieve LDoms Guest Domain Power usage
- 968937 ldm remove-reconf should not take more than 1 guest for argument
- 973950 RFE: Deprecate XML v2 interface in Logical Domains Manager
- 978306 RFE: Direct I/O and IOV support for non-primary Root Domains
- 978990 PM hi water mark must be  $\leq$  DRM's; DRM uses allocated-normalized util, PM use active-normalized util
- 6978992 Binding a domain should cause memory and cpus to occupy the same cpu socket
- 6979530 ldm ls should report normalized utilization information
- 6987355 RFE: LDoms support for Dynamic SR-IOV
- 7022157 Utilization interface should allow averaging history time to be specified
- 7041077 LDoms supports New PM policy that allows only low impact PM states
- 7045735 RFE: Board DR policy engine work for cpu & memory resources
- 7048756 RFE: Dynamic Direct IO (DDIO) support in Logical Domains Manager
- 7056931 "Do not change threading during delayed reconfig" to be explicitly stated in the docs
- 7058702 Primary and guest domain disks corrupted when using a physical disk for guest
- 7061067 RFE: Redesign bootset support
- 7066689 RFE: delayed reconfig enhancements for SR-IOV
- 7079864 RFE: LDoms support for Physical Domain DR
- 7082776 RFE: Need LDoms MIB support on S11
- 7095265 RFE: describe CPU power regions to the affinity engine
- 7099222 RFE: Logical Domains Manager must accept PRI updates for board DR, PM & FMA



- 7107827 Cancelling delayed reconfiguration does not restore extended-mapin-space value.
- 7108333 RFE: Board DR blacklist evacuation support for fault management
- 7117133 Primary panics on repeated delayed reconfig set/cancel of extended-mapin-space parameter
- 7121052 Clean up unneeded support for DR-based power management in code base
- 7121968 Assign reminder or unused interrupt resources to primary domain
- 7125579 ldm returns 0 on failure to start ldg
- 7129548 UUID is not immediately saved in the MD for bound domains if it did not previously exist
- 7130425 Memory leak in parse\_mblock\_str
- 7130693 Whole-core not handled correctly if there is no change to cpus
- 7130880 RFE: rc-add-policy and iov=on/off options to enable IOV features for a Root Complex
- 7130886 RFE: Board DR sequencing of removal/restoring of Root Complexes
- 7130892 RFE: Board DR tracking of evacuated I/O constraints
- 7130897 RFE: Recovery mode support for I/O configuration
- 7131283 Add Oracle VM Server for SPARC 2.2 man pages
- 7131323 Migration of named core domains fails if the target has corresponding core busy or non-existent
- 7131616 RFE: Create Coherency Link nodes in PM and modify PM to report that they are always bound
- 7131701 CPU DR skipped because the last crypto could not be removed from an active domain
- 7133098 Only set-vcpu should work with named cores, plus add|set|rm-core out of sync with add|set|rm-vcpu -c
- 7133885 ldmd dumps core (do\_hv\_pm\_aync) on systems running s10u9\_14a
- 7133898 PM: Add defensive code to guard against invalid SP budget string
- 7134203 SDIO config is not properly rebuilt when 'ldm init-system -r -i filename.xml' is used
- 7138875 P2V not failing on svm soft partitions appropriately
- 7141650 MIB: ldomVdsdevTable is displaying vdsdev details twice
- 7141913 When set-vcpu triggers delayed reconfig, redundant output messages can be seen

- 7142319 Unable to add VFs to the primary domain following the cancellation of delayed reconfiguration
- 7143201 Cross-CPU migration of guests with MAUs to platforms without MAUs fails
- 7143228 RFE: LDoms should retain the whole-core constraint across migrations
- 7143721 Cannot create anymore a Virt. Switch with the Java API, XMPP driver refuses the operation
- 7144570 Migration sometimes failed from a Logical Domains Manager 2.2 system to a 2.1 system
- 7145279 LDoms.2.2\_19 Logical Domains Manager crashes when LDoms 2.1 SDIO config is present
- 7145332 ldm ls-io -p listing has changed in ldoms\_2.2 breaking existing automation
- 7145927 LDoms variables modified by eeprom command at the OS prompt may not persist when ldmd is running
- 7146256 LDoms/SR-IOV: ConnectX device should have better name than "IOVGEN"
- 7146294 LDoms/PM doesn't cope with non-sequential cpu chip IDs, and attempts to park all strands on a node
- 7146725 'ldm init-system -r -i filename.xml' does not reboot primary
- 7146726 Wrong error message for missing vdsdev on target: migration in progress
- 7147314 ldm should not allow empty backend name
- 7148196 Parfait 0.5 catches double free in pm\_pap\_sat.c and multiple file descriptor leaks in various files
- 7148225 Migration of domain with whole-core constraint from 2.2 to 2.1 patch fails
- 7148403 ldm list-io -l -p shows <device>, <device> for buses, rather than <device>, <alias>
- 7149350 RFE: Need to signal to SP when first guest is created by Logical Domains Manager
- 7150238 ldoms-incorporation needs facet.version-lock set in order to allow an unconstrained ldomsmanager
- 7151100 Certain high mblock configs or out of order PA bindings list can lead to assert
- 7151870 Japanese man pages need to be added to the 2.2 IPS manifests
- 7151952 Need to harden XML parser (FUZZ testing)
- 7152464 "Invalid HV MD contents" abort seen on attempting T3-2 power-on with saved max domain SP config
- 7152717 LDoms man pages all refer to "SunOS 5.10" (even on S11)

- 7153153 Box appears to hang when threading is incorrect prior to the whole-core constraint
- 7153753 Affinity bind should call PM function for memory selections
- 7154009 XML representation of guest domain with VF devices does not store VFs as constraints
- 7154012 ldm list-constraints parseable output has 'dev' displaying device alias and is missing device path
- 7154306 'ldm list-spconfig -r' may return "No autosave configurations" when there are valid autosaves
- 7155044 ldmd aborts and restarts on repeated extended-mapin-space set and cancel of del. reconf. operations
- 7155057 quickSuite.py "ldm bind ldg8" fails with "Power Management Failure"
- 7155349 Setting unicast slots to a number exceeding the maximum supported causes it to be reset to 0
- 7155386 ldm init-system command report 'disk server not found' error
- 7156143 ldmd should return failure message to fmd if a page retire request fails
- 7156367 Memory leaks in pm\_testlib
- 7156537 "Solaris\_audit au\_write bad audit descriptor: Bad file number" found during ldoms vcpu setup
- 7156665 ldmd crashed on "fatal error: Frag HV MD assign failed" on powercycle boot to factory-default.
- 7157901 ldmp2v collect phase needs to check zone status more strict
- 7158112 Logical Domains Manager spurious warning: Frag assign for '\_sys\_' /HV MD memory size mismatch, asked for X should be Y
- 7159011 M5/T5: Migration fails initialization on Logical Domains Manager startup
- 7159667 Further vnet addition after system ran out of LDCs caused ldmd SEGV crash and restart
- 7160362 ldmd fails large config spconfig saves due to reduced timeouts
- 7160502 Unable to "re-bind" guest domain if some cpu strands are ILOM disabled
- 7160611 Package system/ldoms/ldomsmanager delivers broken symlink
- 7165101 ldmd coredumps & goes into maintenance mode when attempting rm-io post delayed reconf cancellation
- 7165182 RFE: Add authorization capabilities to the Observability Module

- 7165765 Logical Domains Manager should provide normalized vcpu utilization data for DRM/PM
- 7167990 Delayed reconfig type not persisted across ldmd restarts
- 7170494 RFE: Facility for LDoms to log all guest console messages
- 7171249 ldm ls-io output incorrectly shows niantic PFs even after card is assigned (DIO) to a guest domain
- 7171751 Idle domain with threading=max-ipc shows 88% utilization
- 7171754 Sending XML/XMPP v2 contents but with a v3 version causes ldmd to core dump
- 7172993 RFE: Add to Logical Domains Manager support for platform-specific shared-object modules
- 7173266 Domain with threading=max-ipc loses VCPUs on Logical Domains Manager restart
- 7173298 Logical Domains Manager aborts when it fails to allocate a frag for IO domains
- 7174142 Add pm-address-congruence-offset property to mblock nodes for memory power management
- 7174287 Adding of vnet updates do not generate traps ldomVnetChange trap
- 7174429 Fortify: Replace uses of strcpy with strncpy
- 7175045 ldmd core dumps when boundary case {add|rm}-core operations are attempted
- 7175088 ldomVconsGroupName shows stale information after a change in group name using ldm set-vcons
- 7175500 SNMP MIB functions don't handle new user names in a safe manner
- 7175682 DIOV: remove-io <slot> primary should not dynamically remove slot when using non-diov FW
- 7175721 LDoms 2.2 admin guide: ldomMIB is missing in the mib tree
- 7176144 SNMP MIBs are not described in Oracle VM Server for SPARC 2.2
- 7176492 RFE: LDoms PM must be made a RTL loadable external shared-object
- 7177033 Remove strcpy from mib/\*.c
- 7177928 RFE: Fujitsu M10 Guest Domain error notification feature support on OVM
- 7177931 RFE: Capacity on Demand with OVM for SPARC on Fujitsu M10 system
- 7177934 RFE: HV Debug support
- 7177935 RFE: Support to Collect Information of PCI-Express Card on XSCF

- 7177936 RFE: OVM Fujitsu M10 Power Management support
- 7177938 RFE: Fujitsu M10 RAS feature support on OVM
- 7177940 RFE: Fujitsu M10 OVM supports for SP initiated ordered shutdown
- 7177942 RFE: OVM Fujitsu M10 System Core support
- 7178111 Constraints database needs to encode special chars
- 7178596 set-core fails with a message reporting incorrect number of total cores available
- 7178739 XML list-constraints output needs to encode special chars
- 7178785 install-ldm should print warning and exit if OS is not S10
- 7179408 LDom documentation is not precise in “How to Control Access to All Domain Consoles”
- 7180788 RFE: Create IPS packages for MIB and test libraries
- 7181593 Parfait 1.0 reports File and Memory Leak errors in Logical Domains Manager codebase
- 7181955 RFE: Cross CPU Migration support between Fujitsu M10 server and T series
- 7182069 Change of vnet MTU does not trigger any MIB traps
- 7182744 Change of vsw mtu, vlan and inter-vnet-link failed to trigger traps
- 7182786 LDOMSAG: Chapter 1 Service domain needs minor correction
- 7182849 LDOMSAG: Adding iso image
- 7182850 LDOMSAG: ldm inter-vnet-link property
- 7185428 Device-specific properties of a vf destroyed by board evacuation need to be saved/restored
- 7186029 Adjust PM policy constraint values
- 7186242 ldm ls-devices -a core shows core id out of order after CMU add operations
- 7186417 ldmp2v collect aborts saying patchadd missing instead of saying S11 not supported
- 7186631 Solaris 'eeprom' command timeouts due to MD downloads taking too long
- 7186982 RFE: crash recovery subsystem for Logical Domains Manager
- 7187295 “ldm set-mem --auto-adj 4G primary” failed to adjust memory close to 4G and claim all mem in used
- 7187491 RFE: Add fan and memory power consumption data to Obs Mod
- 7187519 Logical Domains Manager frag code should notify PM when mblock is used/unused

- 7187777 Problem with mgmt/migration
- 7188024 Merge unit test hg repo into Logical Domains Manager source hg repo
- 7188703 Add DBG macros for XML & XMPP and improve general XML debug output
- 7188743 Misc cleanup of XMPP code to harden it against errors
- 7189524 Board DR blacklist pending evacuation completion cleanup for the guest reboot scenario
- 7189675 OPL systems with pkg:/group/system/solaris-large-server not svcs -x clean after installation
- 7190460 ldmd fails: fatal error: xmalloc of negative size (-160) at 274 in ldomcli/vcpu\_cmds.c
- 7190896 DR sequencing fails to cleanup when one of the sequencing event fails
- 7191196 Primary config'd with threading=max-ipc but no whole-core constraint set while in delay-reconf
- 7191523 RFE: list-io command output need to be improved for large systems
- 7191957 ldm add-config -r fails on T4 with SysFW 8.4.0
- 7192032 frag\_pri\_reloc() causing core dump on Fujitsu M10
- 7192369 ldmd core dumps when creating guests (start -i <file>) and failed to create multiple guest domains
- 7193510 pmi\_boottime\_var\_add() does an LDoms database add in delayed reconfig
- 7194230 Explicitly document that setting ldmd/hops to '0' by default does not mean do no checking
- 7195014 Active pm policy not being set transitioning from disabled to elastic policy
- 7195112 Incompatible return type from ldm\_add\_pciepdev()
- 7195243 Repopulation of evacuated IO domains is inefficient
- 7195390 IB SR-IOV: ldmd crashes while creating a VF from an IB SR-IOV Physical function
- 7196466 PM sample\_time calculation may cause overflow
- 7196473 ls-constraints CORE COUNT does not make sense for domains without whole-core
- 7196990 MAC Address collision in a logical domain while migrating or adding new vnet
- 7197678 RFE: Allow multiple vsw per network adaptor on Crossbow to support OVMM live migration

- 7198401 Memory leak in `iov_bind_vfs` (`iov_cmd.c`)
- 7198668 Guest domain panic during `dr_mem?dr_mem_configure+0x94`
- 7198799 Race condition between `handle_mac_response_msg` and `get_remote_ip_addr`
- 7199008 `add-core` and `rm-core` should fail if the domain is not whole-core constrained
- 7199578 Warning: `md_notify_guest_state`: unexpected result 2. Unable to notify SP first guest created
- 7199876 Warning: HV guest start operation failed (7)
- 7199933 RFE: platform information tool for ovs-agent
- 7201510 SNMP OIDs for Fujitsu M10 PM are incorrect
- 7201512 `skip_mblock_t` definition does not match `mblock_t`

