

Oracle® VM Server for SPARC 2.0 Release Notes

Copyright © 2007, 2011, 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 software 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 RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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 which 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.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd.

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.

Copyright © 2007, 2011, Oracle et/ou ses affiliés. Tous droits réservés.

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 RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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

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

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. 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. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Contents

Preface	5
1 Oracle VM Server for SPARC 2.0 Release Notes	9
What's New in This Release	9
System Requirements	10
Supported Platforms	10
Required Software and Patches	11
Related Software	17
Optional Software	17
Software That Can Be Used With the Logical Domains Manager	17
System Controller Software That Interacts With Logical Domains Software	18
Upgrading to Oracle VM Server for SPARC 2.0 Software	18
Upgrade From Logical Domains 1.0 Software Only	19
Known Issues	19
General Issues	19
Bugs Affecting the Oracle VM Server for SPARC 2.0 Software	26
Documentation Errata	49
Resolved Issues	50
Oracle VM Server for SPARC 2.0 RFEs and Bugs Fixed in Oracle Solaris 10 9/10	50
RFEs and Bugs Fixed for Oracle VM Server for SPARC 2.0 Software	52

Preface

The *Oracle VM Server for SPARC 2.0 Release Notes* include 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 Oracle VM Server for SPARC 2.0 software.

Related Documentation

The following table shows the documentation that is available for the Oracle VM Server for SPARC 2.0 release. These documents are available in HTML and PDF formats unless indicated.

TABLE P-1 Related Documentation

Application	Title	Part Number
Oracle VM Server for SPARC 2.0 Software	<i>Oracle VM Server for SPARC 2.0 Administration Guide</i>	821-1485
	<i>Oracle VM Server for SPARC 2.0 Release Notes</i>	821-1487
	<i>Oracle VM Server for SPARC 2.0 Reference Manual</i>	821-1486
	Solaris 10 Reference Manual Collection <ul style="list-style-type: none">▪ <i>drd(1M)</i> man page▪ <i>vntsd(1M)</i> man page	
Oracle VM Server for SPARC Software Basics	<i>Beginners Guide to LDOMs: Understanding and Deploying Logical Domains Software</i> (PDF)	820-0832
Logical Domains Management Information Base (MIB)	<i>Logical Domains (LDOMs) MIB 1.0.1 Administration Guide</i>	820-2319-10
	<i>Logical Domains (LDOMs) MIB 1.0.1 Release Notes</i>	820-2320-10
Oracle Solaris OS: Installation and Configuration	Oracle Solaris 10 9/10 Release and Installation Collection	N/A

You can find documentation that relates to your server, software, or the Oracle Solaris OS on <http://docs.sun.com>. Use the Search box to find the documents and the information that you need.

Documentation, Support, and Training

See information about the following additional resources:

- [Documentation \(http://docs.sun.com\)](http://docs.sun.com)
- [Support \(http://www.sun.com/support/\)](http://www.sun.com/support/)
- [Training \(http://education.oracle.com\)](http://education.oracle.com) – Click the Sun link in the left navigation bar.

Send Oracle Your Comments

Oracle is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to <http://docs.sun.com> and click Feedback.

Include the following book title and part number with your feedback: *Oracle VM Server for SPARC 2.0 Release Notes*, part number 821-1487-12.

Typographic Conventions

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

TABLE P-2 Typographic Conventions

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

Shell Prompts in Command Examples

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

TABLE P-3 Shell Prompts

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

Oracle VM Server for SPARC 2.0 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 2.0 software.

What's New in This Release

The major changes for this release of the Oracle VM Server for SPARC 2.0 software are as follows:

- Support for Oracle's SPARC T3 servers. See “Supported Platforms” on page 10.
- Support for memory dynamic reconfiguration (DR). See “Using Memory Dynamic Reconfiguration” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for the PCI Express (PCIe) direct I/O feature. See “Assigning PCIe Endpoint Devices” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for cooperative guest domain migration. See “Migrating CPUs in an Active Domain” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for CPU whole-core allocation and affinity binding. See “CPU Allocation” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for universally unique identifiers (UUIDs). See “Using Universally Unique Identifiers” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for Logical Domains agents. See the `ldmad(1M)` man page.
- Support for the `virtinfo` command and API. See the `virtinfo(1M)`, `v12n(3EXT)`, and `libv12n(3LIB)` man pages in *Oracle VM Server for SPARC 2.0 Reference Manual*.
- Support for the primary domain being configured from an XML file. See “Saving Domain Configurations for Future Rebuilding” in *Oracle VM Server for SPARC 2.0 Administration Guide*.

- Oracle VM Server for SPARC Physical-to-Virtual (P2V) Conversion Tool enhancements. See [Appendix A, “Oracle VM Server for SPARC Physical-to-Virtual Conversion Tool”](#) in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for memory power management (PM). See “Using Power Management” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for a PM power cap. See “Using Power Management” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for a PM P-state. See “Using Power Management” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Support for virtual disk multipathing. See “Configuring Virtual Disk Multipathing” in *Oracle VM Server for SPARC 2.0 Administration Guide*.
- Bug Fixes

For information about features added to previous versions of the Oracle VM Server for SPARC (Logical Domains) software, see the What's New in Logical Domains Software wiki:

[http://wikis.sun.com/
display/SolarisLogicalDomains/
What%27s+New+in+Oracle+VM+Server+for+SPARC+\(formerly+Logical+Domains\)+Software](http://wikis.sun.com/display/SolarisLogicalDomains/What%27s+New+in+Oracle+VM+Server+for+SPARC+(formerly+Logical+Domains)+Software)

System Requirements

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

Supported Platforms

Note – Starting with the Logical Domains 1.3 release, UltraSPARC T1 platforms are no longer supported. Earlier releases of the Logical Domains software continue to support these platforms.

The Oracle VM Server for SPARC 2.0 software is supported on the following platforms:

- **SPARC T3 servers**
 - SPARC T3-1 server (refer to the *SPARC T3-1 Server Product Notes*)
 - SPARC T3-1B server (refer to the *SPARC T3-1B Server Module Product Notes*)
 - SPARC T3-2 server (refer to the *SPARC T3-2 Server Product Notes*)
 - SPARC T3-4 server (refer to the *SPARC T3-4 Server Product Notes*)
 - Netra SPARC T3-1 server (refer to the *Netra SPARC T3-1 Server Product Notes*)

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

Required Software and Patches

This section lists the required software and patches for use with the Oracle VM Server for SPARC 2.0 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 Logical Domains Manager

After you install the Oracle VM Server for SPARC 2.0 software, you must install patch ID 145880-02 to get all of the bug fixes listed in [“RFEs and Bugs Fixed for Oracle VM Server for SPARC 2.0 Software”](#) on page 52.

To download the patch, see [“Location of Patches”](#) on page 16.

Required and Recommended Oracle Solaris OS

To use all the features of the Oracle VM Server for SPARC 2.0 software, the operating system on all domains should be at least the Oracle Solaris 10 9/10 OS. This OS can be either an initial installation of or an upgrade to the Oracle Solaris 10 9/10 OS.

The following table shows the patches that are required for older releases of the Oracle Solaris 10 OS to have the Oracle VM Server for SPARC 2.0 functionality. These patches are included in the Oracle Solaris 10 9/10 OS.

TABLE 1-1 Patches for Older OS Versions and Domains Needing a Patch

Patch ID	Control Domain	Service Domain	I/O Domain	Guest Domain
141514-02 (vntsd)	X	X		
142909-17 (Oracle Solaris 10 9/10 feature kernel update)	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 2.0 Features

To enable all the Oracle VM Server for SPARC 2.0 features, you must be running at least version 7.3.0 of the system firmware on both UltraSPARC T2 servers and UltraSPARC T2 Plus servers. You must be running at least version 8.0.0 of the system firmware on SPARC T3 servers. This firmware is preinstalled on the SPARC T3 servers. For information about the required Oracle Solaris OS, see [“Required and Recommended Oracle Solaris OS”](#) on page 11.

Required and Recommended System Firmware Patches

To take advantage of all features of Oracle VM Server for SPARC 2.0, ensure that your server runs at least these revisions of the following system firmware patches:

145673-01	Sun SPARC Enterprise T5120 and T5220 Servers
145674-01	Sun Blade T6320 Server Module
145675-01	Netra T5220 Server
145676-01	Sun SPARC Enterprise T5140 and T5240 Servers
145677-01	Netra T5440 Server
145678-01	Sun SPARC Enterprise T5440 Server
145679-01	Sun Blade T6340 Server Module
145680-01	Sun Netra T6340 Server Module
141788-13	Sun Netra CP3260 Blade
139448-12	Sun Netra T6340 Server Module
145665-01	SPARC T3-1 Server

145666-01	SPARC T3-1B Server
145667-01	SPARC T3-2 Server
145668-01	SPARC T3-4 Server
145669-01	Netra SPARC T3-1 Server
145670-01	Netra SPARC T3-1BA Server

Note – SPARC T3 systems are preinstalled with the required firmware.

Minimum Version of Software Required

Running the Oracle VM Server for SPARC 2.0 software along with at least the minimum versions of the other software components specified in [Table 1–2](#) is supported. In such a configuration, you might not get all the features of the Oracle VM Server for SPARC 2.0 software. For production environments, it is best to run with the recommended system firmware version described in [“Required Software to Enable Oracle VM Server for SPARC 2.0 Features” on page 12](#), and with the Oracle Solaris 10 9/10 OS.

An alternate OS strategy is to upgrade the control and service domains to the Oracle Solaris 10 9/10 OS (or to apply the patches listed in [Table 1–1](#)), and to continue running the guest domains at their existing patch level.

Following is a matrix of the minimum versions of software required. The Oracle VM Server for SPARC 2.0 package, SUNWldm, 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).

TABLE 1–2 Minimum Versions of Software

Supported Server	System Firmware	Oracle Solaris OS
SPARC T3 servers	8.0.x	Oracle Solaris 10 9/10 plus the latest revision of patch ID 144488
UltraSPARC T2 Plus servers	7.2.7.x	Solaris 10 8/07 plus patch ID 127111-08 at a minimum
UltraSPARC T2 servers	7.2.7.x	Solaris 10 8/07

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 be running the appropriate software and use supported PCIe cards.

Direct I/O 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.

As of November 2010, the PCIe cards shown in the following table are supported by the direct I/O feature. Refer to your platform's hardware documentation to verify which cards can be used on your platform.

TABLE 1-3 PCIe Cards Supported by the Direct I/O Feature

Device Class	Part Number	PCIe Device Name
Storage	SG-XPCIE2FC-QB4-Z	QLogic 4-Gigabit/Sec PCI-E Dual FC Host Adapter ExpressModule
	SG-XPCIE2FC-QF4	QLogic 4-Gigabit/Sec PCI Express Dual FC Host Adapter
	SG-(X)PCIE1(2)FC-QF8	QLogic 8Gb FC Single and Dual Port
	SG-XPCIE2FC-QF8-Z	QLogic 8-Gigabit/Sec PCI Express Dual FC Host Adapter
	SG-XPCIE8SAS-E-Z	SAS/SATA PCI Express 8-Port Host Adapter
	SGX-SAS6-EM-Z	SAS Eight Port 6-Gigabit/Sec External SAS-2 ExpressModule HBA
	SGX-SAS6-REM-Z	SAS Eight Port 6-Gigabit/Sec SAS-2 REM Internal HBA RAID 0 and 1
	SGX-SAS6-R-REM-Z	Sun Blade 6000 RAID 5 SAS-2 Expansion Module with Write Cache
Network	X1027A-Z	Atlas PCI Express Dual 10-Gigabit Ethernet Fiber
	X4447A-Z	Atlas PCI Express Quad Gigabit Ethernet UTP
	7287A-Z	Atlas Quad Gigabit Ethernet Module
	1108A-Z	Dual 10-Gigabit Ethernet Module
	SX1109A-Z	Dual 10-Gigabit SFP+Ethernet Module
	X4238A-N-Hydra	Hydra (NEM) Pass Through FEM
	X4338A-N Hydra 2	Hydra2 (NEM) Pass Through FEM
	(X)1109a-z	Intel 10 Gbe Fx Dual Port
	(X)1110a-z	Intel 10 Gbe Fx Dual Port
	X7280A-2	Northstar PCI Express Gigabit Ethernet UTP

TABLE 1-3 PCIe Cards Supported by the Direct I/O Feature (Continued)

Device Class	Part Number	PCIe Device Name
	7282A-Z	Northstar Dual Gigabit Ethernet UTP Module
	X2073A-N	Opus Intel PCIe2 10 Gbe FEM
	X7283A-Z	Sun PCI-E Express Dual Port Gigabit Ethernet MMF (fibre) ExpressModule
	SESX7XA1Z	XAUI 10-Gigabit Ethernet

Direct I/O Software Requirements

To use the DIO feature, the following domains must be running the supported OS:

- **primary domain.** At least the Oracle Solaris 10 9/10 OS plus patch ID 145868-01
- **I/O domain.** Any Oracle Solaris OS that is supported by the platform

Location of Oracle VM Server for SPARC 2.0 Software

You can find the Oracle VM Server for SPARC 2.0 software to download at <http://www.oracle.com/virtualization/index.html>.

Note – The Oracle VM Server for SPARC 2.0 software no longer comes bundled with Solaris Security Toolkit (SST) to provide security services. For more information about SST, see “Optional Software” on page 17.

The OVM_Server_SPARC-2_0.zip file that you download contains the following:

- Oracle VM Server for SPARC 2.0 software (SUNWldm.v)
- The ld(1M), ldmconfig(1M), and ldmd(1M) man pages in the SUNWldm.v package are installed when the package is installed
- Installation script for Oracle VM Server for SPARC 2.0 software (install-ldm)
- Logical Domains Management Information Base (SUNWldmib.v)
- Physical-to-Virtual Conversion Tool (SUNWldmp2v)
- Configuration Assistant GUI (Configurator.jar)

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

```
OVM_Server_SPARC-2_0/
  Install/
    install-ldm
  Product/
    Configurator/
      Configurator.jar
```

README.GUI
SUNWldm.v
SUNWldmib.v
SUNWldmp2v
README

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 Logical Domains Manager and Oracle Solaris OS patches at <http://sunsolve.sun.com> or at <http://support.oracle.com>.

Location of Documentation

The *Oracle VM Server for SPARC 2.0 Administration Guide*, the *Oracle VM Server for SPARC 2.0 Reference Manual*, and these *Oracle VM Server for SPARC 2.0 Release Notes* can be obtained from:

<http://docs.sun.com/app/docs/prod/ldoms>

The Oracle VM Server for SPARC 2.0 man pages are installed on your system as part of the SUNWldm.v and SUNWldmp2v packages.

The Sun Logical Domains Wiki contains Best Practices, Guidelines, and Recommendations for deploying Logical Domains software.

<http://wikis.sun.com/display/SolarisLogicalDomains/Home>

The *Beginners Guide to LDomS: Understanding and Deploying Logical Domains* can be used to get a general overview of Logical Domains software. However, the details of the guide specifically apply to the Logical Domains 1.0 software release and are now out of date for Oracle VM Server for SPARC 2.0 software. The guide can be found at the Sun BluePrints site.

<http://www.sun.com/blueprints/0207/820-0832.html>

Related Software

Optional Software

This section describes software that you can use with the Logical Domains software.

- **Solaris Security Toolkit (SST) software** can help you secure the Oracle Solaris OS in the control domain and other domains. SST is no longer supported and is no longer bundled with the Oracle VM Server for SPARC 2.0 software.

You can still download and use the SST 4.2 to harden your systems. SST 4.2, which was previously bundled with Logical Domains releases, is still available from <http://www.sun.com/software/security/jass/>.

- **Logical Domains Management Information Base (MIB) software** can help you enable third-party applications to perform remote monitoring and a few control operations. Refer to the *Logical Domains (LDoms) MIB 1.0.1 Administration Guide* and the *Logical Domains (LDoms) MIB 1.0.1 Release Notes* for more information.

Software That Can Be Used With the Logical Domains Manager

This section details the software that is compatible with and can be used with the Logical Domains software. Be sure to check the software documentation or your platform documentation to find the version number of the software that is available for your version of the Logical Domains software and your platform.

- **SunVTS** functionality is available in the control domain and guest domains on certain Logical Domains software releases and certain platforms. SunVTS is a validation test suite that provides a comprehensive diagnostic tool that tests and validates Oracle's Sun hardware by verifying the connectivity and proper functioning of most hardware controllers and devices on Oracle's Sun servers. For more information about SunVTS, refer to the *SunVTS 7.0 Software* (<http://dlc.sun.com/pdf/821-1360-12/821-1360-12.pdf>).
- **Explorer Data Collector** can be used with the Logical Domains Manager software enabled on the control domain. Explorer is a diagnostic data collection tool. The tool comprises shell scripts and a few binary executables. For more information, see the *Oracle Explorer User's Guide*.
- **Solaris Cluster** software can be used only on an I/O domain in Logical Domains software releases up through Logical Domains 1.0.2. In Logical Domains 1.0.3, 1.1, 1.2, and 1.3 software, Solaris Cluster software can be used in a guest domain with some restrictions. Refer to Solaris Cluster documentation for more information about any restrictions and

about the Solaris Cluster software in general. Starting with the Logical Domains 1.2 release and the Solaris Cluster 11/09 release, you can also manage logical domains as a resource by using the 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 Solaris Containers and Logical Domains, and supports hardware management from power up to production. For more information, see <http://www.oracle.com/us/products/enterprise-manager/opscenter/index.html>.

System Controller Software That Interacts With Logical Domains Software

The following system controller (SC) software interacts with the Oracle VM Server for SPARC 2.0 software:

- **Sun Integrated Lights Out Manager (ILOM) 3.0** is the system management firmware that you can use to monitor, manage, and configure UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 server platforms. ILOM is preinstalled on these platforms and can be used on supported servers with the Oracle VM Server for SPARC 2.0 software enabled. Refer to the *Sun Integrated Lights Out Manager 3.0 User's Guide* for features and tasks that are common to Oracle's Sun rackmounted servers or blade servers that support ILOM. Other user documents present ILOM features and tasks that are specific to the server platform that you are using. You can find the ILOM platform-specific information within the documentation set that accompanies your system.
- **Netra Data Plane Software Suite** is a complete board software package solution. The software provides an optimized rapid development and runtime environment on top of multistrand partitioning firmware for Sun CMT platforms. The Logical Domains Manager contains some `ldm` subcommands (`add-vdpcs`, `rm-vdpcs`, `add-vdpcc`, and `rm-vdpcc`) for use with this software. Refer to the *Netra Data Plane Software Suite 2.0 User's Guide* for more information about this software.

Upgrading to Oracle VM Server for SPARC 2.0 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 2.0 software, see [“Upgrade to the Oracle VM Server for SPARC 2.0 Software”](#) in *Oracle VM Server for SPARC 2.0 Administration Guide*.

Upgrade From Logical Domains 1.0 Software Only

Note – You *cannot* upgrade to the Oracle VM Server for SPARC 2.0 software from an UltraSPARC T1 system that runs the Logical Domains 1.0 software. The UltraSPARC T1 platform is *not* supported by the Oracle VM Server for SPARC 2.0 software.

Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 2.0 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.

I/O MMU Bypass Mode Is No Longer Needed

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

Service Processor and System Controller Are Interchangeable Terms

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

Cards Not Supported

The following cards are *not* supported for this Oracle VM Server for SPARC 2.0 software release:

- Sun Dual Port 4x IB Host Channel Adapter PCI-X Card
- Dual Port 4x PCI EXPRESS Infiniband Host Channel Adapter - Low Profile



Caution – If these unsupported configurations are used with Oracle VM Server for SPARC 2.0, stop and unbind all logical domains before the control domain is rebooted. Failure to do so can result in a system crash causing the loss of all the logical domains that are active in the system.

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 9/10 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 9/10, the physical disk slice exported as a virtual disk will appear in the guest domain with no device ID.

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

Workaround: After upgrading a service domain to Oracle Solaris 10 9/10, if a guest domain is unable to find its Solaris Volume Manager configuration or its metadevices, execute the following procedure.

▼ 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 dev_id for 'vdc', 0x10
```

These messages are normal and do not report any problems.

Logical Domain Channels and Logical Domains

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

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

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

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

1. The control domain allocates 12 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.
2. The control domain allocates 1 LDC to every logical domain, including itself, for control traffic.
3. Each virtual I/O service on the control domain consumes 1 LDC for every connected client of that service.

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

- Virtual network
- Virtual disk
- Virtual console

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

$12(1) + 9(2) + 8 \times 3(3) = 45$ 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:

$12 + 46 + 45 \times 11 = 553$ LDCs in total.

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

Memory Size Requirements

The Oracle VM Server for SPARC software does not impose a memory size limitation when you create a domain. The memory size requirement is a characteristic of the guest operating system. Some Oracle VM Server for SPARC functionality might not work if the amount of memory present is less than the recommended size. For recommended and minimum size memory requirements for the Oracle Solaris 10 OS, see [“System Requirements and Recommendations” in Oracle Solaris 10 9/10 Installation Guide: Planning for Installation and Upgrade](#).

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

The memory dynamic reconfiguration (DR) feature enforces 256-Mbyte alignment on the address and size of the memory involved in a given operation. See [“Memory Alignment” in Oracle VM Server for SPARC 2.0 Administration Guide](#).

Booting a Large Number of Domains

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

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

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

To run the maximum configurations, a machine needs the following amount of memory, depending on your platform, so that the guest domains contain an adequate amount of memory:

- 128 Gbytes of memory for SPARC T3 servers
- 128 Gbytes of memory for UltraSPARC T2 Plus servers
- 64 Gbytes of memory for UltraSPARC T2 servers

Memory and swap space usage increases in a guest domain when the vsw services used by the domain provides services to many virtual networks (in multiple domains). This is due to the peer-to-peer links between all the vnet connected to the vsw. The service domain benefits from

having extra memory. Four Gbytes is the recommended minimum when running more than 64 domains. Start domains in groups of 10 or less and wait for them to boot before starting the next batch. The same advice applies to installing operating systems on domains.

Cleanly Shutting Down and Power Cycling a Logical Domains System

If you have made any configuration changes since last saving a configuration to the SC, before you attempt to power off or power cycle a Logical Domains system, make sure that you save the latest configuration that you want to keep.

▼ Power Off a System With Multiple Active Domains

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

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

▼ Power Cycle the System

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

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

Memory Size Requested Might Be Different From Memory Allocated

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

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

Logical Domains Variable Persistence

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

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

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

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

- At the OpenBoot prompt
- Using the Oracle Solaris OS `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 2.0 software, there are a few cases where variable updates do not persist as expected:

- All methods of updating a variable persist across reboots of that domain. However, they do not persist across a powercycle of the system, unless a subsequent logical domain configuration is saved to the SC. The methods of updating a variable include by OpenBoot firmware and by the `eeeprom` and `ldm` commands. In addition, in the control domain, updates made using OpenBoot firmware persist across a powercycle of the system, that is, even without subsequently saving a new logical domain configuration to the SC.
- In all cases, when reverting to the factory-default configuration from a configuration generated by the Logical Domains Manager, all Logical Domains variables start with their default values.

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

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


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

- When running Live Upgrade, perform the following steps:

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

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

The following **Bug IDs** have been filed to resolve these issues: **6520041**, **6540368**, **6540937**, and **6590259**.

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

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

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

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

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

Workaround: Set the performance mode for the power management policy.

Fault Management

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

Delayed Reconfiguration

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

Cryptographic Units

The Oracle Solaris 10 10/09 OS introduces the capability to dynamically add and remove cryptographic units from a domain, which is called cryptographic unit dynamic reconfiguration (DR). The Logical Domains Manager automatically detects whether a domain allows cryptographic unit DR, and enables the functionality only for those domains. In addition, CPU DR is no longer disabled in domains that have cryptographic units bound and are running an appropriate version of the Oracle Solaris OS.

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

Bugs Affecting the Oracle VM Server for SPARC 2.0 Software

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

Logical Domains Manager Does Not Validate Disk Paths and Network Devices

Bug ID 6447740: The Logical Domains Manager does not validate disk paths and network devices.

Disk Paths

If a disk device listed in a guest domain's configuration is either non-existent or otherwise unusable, the disk cannot be used by the virtual disk server (vds). However, the Logical Domains Manager does not emit any warning or error when the domain is bound or started.

When the guest tries to boot, messages similar to the following are printed on the guest's console:

```
WARNING: /virtual-devices@100/channel-devices@200/disk@0: Timeout
connecting to virtual disk server... retrying
```

In addition, if a network interface specified using the `net - dev=` property does not exist or is otherwise unusable, the virtual switch is unable to communicate outside the physical machine, but the Logical Domains Manager does not emit any warning or error when the domain is bound or started.

▼ Recover From an Errant `net - dev` Property Specified for a Virtual Switch

- 1 Issue the `ldm set -vsw` command with the corrected `net - dev` property value.
- 2 Reboot the domain hosting the virtual switch in question.

▼ Recover From an Errant Virtual Disk Service Device or Volume

- 1 Stop the domain owning the virtual disk bound to the errant device or volume.
- 2 Issue the `ldm rm -vdsdev` command to remove the errant virtual disk service device.
- 3 Issue the `ldm add -vdsdev` command to correct the physical path to the volume.
- 4 Restart the domain owning the virtual disk.

Network Devices

If a disk device listed in a guest domain's configuration is being used by software other than the Logical Domains Manager (for example, if it is mounted in the service domain), the disk cannot be used by the virtual disk server (vds), but the Logical Domains Manager does not emit a warning that it is in use when the domain is bound or started.

When the guest domain tries to boot, a message similar to the following is printed on the guest's console:

```
WARNING: /virtual-devices@100/channel-devices@200/disk@0: Timeout  
connecting to virtual disk server... retrying
```

▼ Recover From a Disk Device Being Used by Other Software

- 1 Unbind the guest domain.
- 2 Unmount the disk device to make it available.
- 3 Bind the guest domain.
- 4 Boot the domain.

Hang Can Occur With Guest OS in Simultaneous Operations

Bug ID 6497796: Under rare circumstances, when a Logical Domains variable, such as `boot-device`, is being updated from within a guest domain by using the `eeprom(1M)` command at the same time that the Logical Domains Manager is being used to add or remove virtual CPUs from the same domain, the guest OS can hang.

Workaround: Ensure that these two operations are not performed simultaneously.

Recovery: Use the `ldm stop-domain` and `ldm start-domain` commands to stop and start the guest OS.

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

Bug ID 6506494: There are some cases where the behavior of the `ldm stop-domain` command is confusing.

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

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

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

Cannot Set Security Keys With Logical Domains Running

Bug ID 6510214: In a Logical Domains environment, there is no support for setting or deleting wide-area network (WAN) boot keys from within the Oracle Solaris OS by using the `ickey(1M)` command. All `ickey` operations fail with the following error:

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

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

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

Bug ID 6540368: This issue is summarized in “[Logical Domains Variable Persistence](#)” on [page 24](#) and affects only the control domain.

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

Bug ID 6591844: If a CPU or memory fault occurs, the affected domain might panic and reboot. If the Fault Management Architecture (FMA) attempts to retire the faulted component

while the domain is rebooting, the Logical Domains Manager is not able to communicate with the domain, and the retire fails. In this case, the `fmadm faulty` command lists the resource as degraded.

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

```
primary# svcadm restart fmd
```

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

Bug ID 6603974: If you configure more than four virtual networks (`vnets`) in a guest domain on the same network using the Dynamic Host Configuration Protocol (DHCP), the guest domain can eventually become unresponsive while running network traffic.

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

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

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

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

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

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

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

Bug ID 6656033: Simultaneous net installation of multiple guest domains fails on Sun SPARC Enterprise T5140 and Sun SPARC Enterprise T5240 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.

Solaris Volume Manager Volumes Built on Slice 2 Fail JumpStart When Used as the Boot Device in a Guest Domain

Bug ID 6687634: If the Solaris Volume Manager volume is built on top of a disk slice that contains block 0 of the disk, then Solaris Volume Manager prevents writing to block 0 of the volume to avoid overwriting the label of the disk.

If an Solaris Volume Manager volume built on top of a disk slice that contains block 0 of the disk is exported as a full virtual disk, then a guest domain is unable to write a disk label for that virtual disk, and this prevents the Oracle Solaris OS from being installed on such a disk.

Workaround: Solaris Volume Manager volumes exported as a virtual disk should not be built on top of a disk slice that contains block 0 of the disk.

A more generic guideline is that slices that start on the first block (block 0) of a physical disk should not be exported (either directly or indirectly) as a virtual disk. Refer to [“Directly or Indirectly Exporting a Disk Slice”](#) in *Oracle VM Server for SPARC 2.0 Administration Guide*.

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

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

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

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

Bug ID 6742805: A domain shutdown or memory scrub can take over 15 minutes with a single CPU and a very large memory configuration. During a shutdown, the CPUs in a domain are used to scrub all the memory owned by the domain. The time taken to complete the scrub can be quite long if a configuration is imbalanced, for example, a single CPU domain with 512 Gbytes of memory. This prolonged scrub time extends the amount of time it takes to shut down a domain.

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

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

Bug ID 6753683: Sometimes, executing the `uadmin 1 0` command from the command line of an Logical Domains system does not leave the system at the ok prompt after the subsequent reset. This incorrect behavior is seen only when the Logical Domains variable `auto-reboot?` is set to `true`. If `auto-reboot?` is set to `false`, the expected behavior occurs.

Workaround: Use this command instead:

```
uadmin 2 0
```

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

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

Bug ID 6760933: On occasion, an active logical domain appears to be in the *transition* state instead of the *normal* state long after it is booted or following the completion of a domain migration. This glitch is harmless, and the domain is fully operational. To see what flag is set, check the `flags` field in the `ldm list -l -p` command output, or check the `FLAGS` field in the `ldm list` command, which shows `-n----` for *normal* or `-t----` for *transition*.

Recovery: The logical domain should display the correct state upon the next reboot.

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

Bug ID 6764613: If you do not have a network configured on your machine and have a Network Information Services (NIS) client running, the Logical Domains Manager will not start on your system.

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

```
# svcadm disable nis/client
```

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

Bug ID 6772089: In certain situations, a migration fails and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This can occur even if the total amount of available memory on the target machine is greater than the amount of memory being used by the source domain.

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

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

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

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

Bug ID 6772120: If the virtual disk on the target machine does not point to the same disk back end that is used on the source machine, the migrated domain cannot access the virtual disk using that disk back end. A hang can result when accessing the virtual disk on the domain.

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

Workaround: Ensure that when you are configuring the target domain to receive a migrated domain that the disk volume (`vdsdev`) matches the disk back end used on the source domain.

Recovery: Do one of the following if you discover that the virtual disk device on the target machine points to the incorrect disk back end:

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

Constraint Database Is Not Synchronized to Saved Configuration

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

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

Explicit Console Group and Port Bindings Are Not Migrated

Bug ID 6781589: During a migration, any explicitly assigned console group and port are ignored, and a console with default properties is created for the target domain. This console is created using the target domain name as the console group and using any available port on the first virtual console concentrator (vcc) device in the control domain. If there is a conflict with the default group name, the migration fails.

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

VIO DR Operations Ignore the Force (-f) Option

Bug ID 6703127: Virtual input/output (VIO) dynamic reconfiguration (DR) operations ignore the `-f` (force) option in CLI commands.

ldmconfig Might Cause the Root File System of the Control Domain to Become Full and Halt the System

Bug ID 6848114: `ldmconfig` can run on a system that does not have file systems of sufficient capacity to contain the virtual disks for the created domains. In this situation, an error message is issued. However, `ldmconfig` permits you to continue to use the disks that are in `/ldoms/disks` to deploy the configuration. This situation could cause the root file system of the control domain to become full and halt the system.

Workaround: Do the following:

1. Exit the Configuration Assistant by typing `q` or by typing `Ctrl-C`.
2. Add more file systems of adequate capacity.
3. Rerun the `ldmconfig` command.

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

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

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

Workaround: Reboot the guest domain.

ldmd Might Dump Core If Multiple set - vcpu Operations Are Performed on the Control Domain While It Is in Delayed Reconfiguration Mode

Bug ID 6697096: Under certain circumstances, when multiple `ldm set -vcpu` operations are performed on the control domain while it is in delayed reconfiguration mode, `ldmd` might abort and be restarted by the Service Management Facility (SMF).

While the control domain is in delayed reconfiguration mode, take care when attempting an `ldm set -vcpu` operation. A single `ldm set -vcpu` operation will succeed, but a second `ldm set -vcpu` operation might cause the `ldmd` daemon to dump core.

Workaround: Reboot the control domain before you attempt the second `ldm set -vcpu` operation.

Virtual Network Devices Are Not Created Properly on the Control Domain

Bug ID 6836587: Sometimes `ifconfig` indicates that the device does not exist after you add a virtual network or virtual disk device to a domain. This situation might occur as the result of the `/devices` entry not being created.

Although this should not occur during normal operation, the error was seen when the instance number of a virtual network device did not match the instance number listed in `/etc/path_to_inst` file.

For example:

```
# ifconfig vnet0 plumb
ifconfig: plumb: vnet0: no such interface
```

The instance number of a virtual device is shown under the `DEVICE` column in the `ldm list` output:

```
# ldm list -o network primary
NAME
primary

MAC
00:14:4f:86:6a:64

VSW
NAME          MAC          NET-DEV DEVICE  DEFAULT-VLAN-ID PVID VID  MTU  MODE
primary-vsw0 00:14:4f:f9:86:f3 nxge0  switch@0  1              1    1500

NETWORK
NAME  SERVICE          DEVICE  MAC          MODE PVID VID  MTU
vnet1 primary-vsw0@primary network@0 00:14:4f:f8:76:6d  1    1500
```

The instance number (0 for both the `vnet` and `vsw` shown previously) can be compared with the instance number in the `path_to_inst` file to ensure that they match.

```
# egrep '(vnet|vsw)' /etc/path_to_inst
"/virtual-devices@100/channel-devices@200/virtual-network-switch@0" 0 "vsw"
"/virtual-devices@100/channel-devices@200/network@0" 0 "vnet"
```

Workaround: In the case of mismatching instance numbers, remove the virtual network or virtual switch device. Then, add them again by explicitly specifying the instance number required by setting the `id` property.

You can also manually edit the `/etc/path_to_inst` file. See the [path_to_inst\(4\)](#) man page.



Caution – Be aware of the warning contained in the man page that states “changes should not be made to `/etc/path_to_inst` without careful consideration.”

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

Bug ID 6757486: Occasionally, after a domain has been migrated, it is not possible to connect to the console for that domain.

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

```
# svcadm restart vntsd
```

Note – This command will disconnect all active console connections.

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

Bug ID 6808832: You can configure a maximum of two domains with dedicated PCI-E root complexes on systems such as the Sun Fire T5240. These systems have two UltraSPARC T2+ CPUs and two I/O root complexes.

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

The `pci@400` fabric (or leaf) contains the onboard `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.

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

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

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

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

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

ldm Commands Are Slow to Respond When Several Domains Are Booting

Bug ID 6855079: An `ldm` command might be slow to respond when several domains are booting. If you issue an `ldm` command at this stage, the command might appear to hang. Note that the `ldm` command will return after performing the expected task. After the command returns, the system should respond normally to `ldm` commands.

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

Spurious ds_ldc_cb: LDC READ event Message Seen When Rebooting the Control Domain or a Guest Domain

Bug ID 6846889: When rebooting the control domain or a guest domain, the following warning message might be logged on the control domain and on the guest domain that is rebooting:

```
WARNING: ds@0: ds_ldc_cb: LDC READ event while port not up
```

Workaround: You can ignore this message.

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

Bug ID 6904849: Starting with the Logical Domains 1.3 release, a domain can be migrated even if it has one or more cryptographic units bound to it.

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

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

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

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

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

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

Confusing Migration Failure Message for Real Address Memory Bind Failures

Bug ID 6904240: In certain situations, a migration fails with the following error message, and `ldmd` reports that it was not possible to bind the memory needed for the source domain. This situation can occur even if the total amount of available memory on the target machine is greater than the amount of memory being used by the source domain (as shown by `ldm ls-devices -a mem`).

Unable to bind 29952M memory region at real address 0x8000000
Domain Migration of LDom ldg0 failed

Cause: This failure is due the inability to meet congruence requirements between the Real Address (RA) and the Physical Address (PA) on the target machine.

Workaround: Stop the domain and perform the migration as a cold migration. You can also reduce the size of the memory on the guest domain by 128 Mbytes, which might permit the migration to proceed while the domain is running.

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

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

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

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

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

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

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

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

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

3. **Perform a reconfiguration reboot of the control domain.**

```
# reboot -- -r
```

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

Dynamically Removing All the Cryptographic Units From a Domain Causes SSH to Terminate

Bug ID 6897743: If all the hardware cryptographic units are dynamically removed from a running domain, the cryptographic framework fails to seamlessly switch to the software cryptographic providers, and kills all the ssh connections.

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

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

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

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

Bug ID 6892229: When you run the `ldm ls-io -l` command on a system that has an Atlas 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 set-domain -i` Command Fails to Completely Configure the Domain as Specified by the XML File

Bug ID 6965623: The `ldm set-domain -i filename` command cannot create a logical domain based on the configuration described by the specified XML file. The command parses the `ldom_info` node from the file, but fails to parse the other resource nodes, such as `cpu`, `mau`, and `memory`.

Workaround: Use the `ldm add-domain -i filename` command to add the domain and to configure the domain based on the specified XML file.

`ldmp2v` Command: `ufsdump` Archiving Method Is No Longer Used

Bug ID 6924550: When you use the `ufsdump` archiving method and a virtual disk backed by a file on a UFS file system, the system might hang when you run the `ldmp2v prepare` command. The `ufsdump` archiving method is no longer available for the `ldmp2v collect` command and has been superseded by the Flash archiving method. For compatibility with previously created `ufsdump` archives, the `ldmp2v prepare` command can still use existing `ufsdump` archives. However, it is better to create a new archive by using the Flash archiving method. This potential hang also applies 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.

ldmp2v convert Command: VxVM Warning Messages During Boot

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

```

Boot device: disk0:a File and args:
SunOS Release 5.10 Version Generic_139555-08 64-bit
Copyright 1983-2009 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Hostname: normal
Configuring devices.
/kernel/drv/sparcv9/vxdmp: undefined symbol ?romp?
WARNING: mod_load: cannot load module ?vxdmp?
WARNING: vxdmp: unable to resolve dependency, module ?misc/ted? not found
/kernel/drv/sparcv9/vxdmp: undefined symbol ?romp?
WARNING: mod_load: cannot load module ?vxdmp?
WARNING: vxdmp: unable to resolve dependency, module ?misc/ted? not found
/kernel/drv/sparcv9/vxio: undefined symbol ?romp?
WARNING: mod_load: cannot load module ?vxio?
WARNING: vxio: unable to resolve dependency, module ?drv/vxdmp? not found
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
NOTICE: VxVM not started

```

Dynamic Reconfiguration of MTU Values of Virtual Network Devices Sometimes Fails

Bug ID 6936833: If you modify the maximum transmission unit (MTU) of a virtual network device on the control domain, a delayed reconfiguration operation is triggered. If you subsequently cancel the delayed reconfiguration, the MTU value for the device is not restored to the original value.

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

```

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

```


Active Guest Domain Erroneously Shown to Be in Transition Mode

Bug ID 6961910: An active guest domain that runs the Oracle Solaris OS appears to be in the transition state rather than in the normal state in the following circumstances:

- The CPU with the lowest physical ID is dynamically removed from the domain.
- Subsequently, a new CPU is dynamically added.

Workaround: Reboot the domain to return it to the normal state.

ldmconfig -c Command Fails Silently

Bug ID 6922142: When you run the `ldmconfig -c` command to check the Oracle Solaris OS media for valid packages, the check fails silently. The appropriate virtual devices to access the CD-ROM are not created, so the newly created logical domains cannot boot from or install from a CD-ROM device.

The `ldmconfig` command does not work if a user installs the Oracle Solaris OS from a CD.

Workaround: If you boot from a CD, you must manually configure the logical domains. Manually create the CD-ROM `cdsdev` and virtual disk. For example, for the Oracle Solaris 10/09 OS and the `ldg0` domain, you would type the following commands:

```
# ldm add-vdsdev options=ro /vol/dev/dsk/c0t0d0/sol_10_1009_sparc cdrom0@primary-vds0
# ldm add-vdisk cdrom cdrom0@primary-vds0 ldg0
```

Memory DR Is Disabled Following a Canceled Migration

Bug ID 6956431: After an Oracle Solaris 10 9/10 domain has been suspended as part of a migration operation, memory dynamic reconfiguration (DR) is disabled. This applies not only when the migration is successful, but also when the migration has been canceled, despite the fact that the domain remains on the source machine.

Migration of a Guest Domain That Has Hybrid I/O-Enabled Virtual Network Devices Panics the Service Domain

Bug ID 6972633: The service domain panics when performing a warm migration of a guest domain. The source system in the migration is a SPARC T3-1 that has the NIU hybrid I/O capability.

The problem can occur when *all* of the following conditions are met:

- The source system is a SPARC T3-1 server.
- The guest domain has hybrid I/O mode enabled.
- The guest domain has a hybrid I/O resource assigned.

A guest domain that has hybrid I/O enabled for a virtual network interface shows `hybrid` in the `MODE` column as follows:

```
# ldm list -o network ldg1
...
NAME      SERVICE          ID DEVICE      MAC              MODE  PVID  MTU
vnet2     niu-vsw@primary  1  network@1  00:14:4f:fa:9e:89  hybrid  1    1500
```

However, the hybrid I/O resource is only assigned if the following command shows any output on the guest domain:

```
# kstat -p nxge
```

Workaround: Perform the following steps:

1. Obtain the current configuration of the virtual network device.

This step ensures that replumbing the interface is error-free.

```
# ifconfig vnet1
```

2. Unplumb the virtual network interface on the guest domain prior to the migration.

```
# ifconfig vnet1 unplumb
```

3. Perform the migration.

4. Plumb the interface.

```
# ifconfig vnet1 plumb
```

Migrated Guest Domain Hangs If It Has CPUs That Are Offline, Faulted, or Spare

Bug ID 6965789: This hang occurs on a guest domain under the following conditions:

- The domain is running the Oracle Solaris 10 9/10 OS.
- The source and target machines are running the Oracle VM Server for SPARC 2.0 release.
- The domain has been migrated by using warm migration.
- The domain has CPUs that have been marked as offline, faulted, or spare.

Use the `ps rinfo` command to determine whether any CPUs in the guest domain have been marked as offline, faulted, or spare. This command shows the state of CPUs.

Workaround: To avoid the problem on a guest domain that has offline or spare CPUs, use the `ps radm` command to bring all CPUs online before attempting to migrate the domain.

ldmconfig Command Does Not Save Virtual Disk Service Devices for Bound Domains

Bug ID 6967799: When you use the `ldmconfig` command to create a service processor (SP) configuration, the default `primary-with-clients` configuration is created before the command adds the `vol0` or `vol1` virtual disk service device.

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

Bug ID 6968507: Following the migration of an active domain, CPU utilization in the migrated domain can increase dramatically for a short period of time. If there is a dynamic resource management (DRM) policy in effect for the domain at the time of the migration, the Logical Domains Manager might begin to add CPUs. In particular, if the `vcpu-max` and `at tack` properties were not specified when the policy was added, the default value of `unlimited` causes all the unbound CPUs in the target machine to be added to the migrated domain.

Recovery: No recovery is necessary. After the CPU utilization drops below the upper limit that is specified by the DRM policy, the Logical Domains Manager automatically removes the CPUs.

Uncooperative Oracle Solaris Domain Migration Can Be Blocked If `cpu0` Is Offline

Bug ID 6965758: The migration of an active domain can fail if it is running a release older than the Oracle Solaris 10 10/09 OS and the lowest numbered CPU in the domain is in the `offline` state. The operation fails when the Logical Domains Manager uses CPU DR to reduce the domain to a single CPU. In doing so, the Logical Domains Manager attempts to remove all but the lowest CPU in the domain, but as that CPU is offline, the operation fails.

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

Sometimes DR Requests Fail to Remove All Requested CPUs

Bug ID 6493140: Sometimes, the Oracle Solaris OS is unable to use DR to remove all the requested CPUs. When this problem occurs, you see error messages similar to the following:

```
Removal of cpu 10 failed
```

Recovery: Issue a subsequent request to remove the number of CPUs that failed to be removed the first time. Such a retry generally succeeds.

When Re-creating a Guest Domain From an XML File, the Domain Should be Created by Using the CPU Allocation Unit and Not the Whole Core Number

Bug ID 6977065: A guest domain that is re-created by means of an XML file incorrectly uses the value of the `max-cores` property instead of the CPU allocation unit.

For example, a guest domain that is initially configured and bound with 10 whole cores (`max-cores=10`) is later reconfigured to use three whole cores. At this time, the guest domain's constraints are saved to an XML file. Later, the guest domain is unbound and destroyed, and

then re-created by using the saved XML file. Instead of the guest domain being configured with three whole cores, it is configured with 10 whole cores based on the value of the `max-cores` property.

Inactive or Bound Domains Allow the Coexistence of the Whole-Core Constraint With an Enabled DRM Policy

Bug ID 6978606: The whole-core constraint and a DRM policy are not permitted to be enabled simultaneously. However, depending on the order in which you enable the constraint or policy, the `ldm list` command shows that both are enabled simultaneously.

If you enable the whole-core constraint on an inactive or bound domain that already has an active DRM policy, both appear as enabled in the `ldm list` output. In this situation, the DRM policy is disabled when you make the domain active.

Ldm set-vcc port-range Command While in Delayed Reconfiguration Mode Fails to Hold the New port-range Value After Reboot

Bug ID 6975322: When the primary domain is in delayed reconfiguration mode, the `ldm set-vcc port-range` command does not work.

Workaround: Ensure that the primary domain is not in delayed reconfiguration mode before running the `ldm set-vcc port-range` command.

Using ldm stop -a Command on Domains in a Master-Slave Relationship Leaves the Slave With the stopping Flag Set

Bug ID 6979574: When a reset dependency is created, an `ldm stop -a` command might result in a domain with a reset dependency being restarted instead of only stopped.

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

Guest Domain Might Panic After the Domain Is Started With the ldm start Command

Bug ID 6979007: If you attempt to start a guest domain in elastic mode with the `ldm start` command, the guest domain might panic during boot.

This panic is rare and can only be seen from the guest domain's console. The control domain is not affected, and there is no indication that the guest domain panicked on the control domain console.

Workaround: No action is necessary. The guest domain resets and reboots successfully.

ldmp2v convert Command: The drd Daemon Is Not Always Enabled

Bug ID 6924931: When you convert a sun4u system that runs the Oracle Solaris 10 OS to a logical domain, the drd daemon is not enabled automatically after conversion. As a result, DR is not possible for the converted domain.

Workaround: Manually enable the drd daemon after the conversion by running the following command:

```
# svcadm enable svc:/platform/sun4v/drd:default
```

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

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

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

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

Workaround: You can ignore these messages.

Memory DR Removal Operations With Unplumbed NIU nxge Instances Can Cause a Kernel Panic, Panic - Generated at user request

Bug ID 6977504: When a memory dynamic reconfiguration (DR) removal command (`ldm rm-mem` or `ldm set-mem`) is issued that affects a domain that owns an NIU device, the domain might panic if the NIU's nxge driver instance is not plumbed. The panic might look similar to the following:

```
panic[cpu0]/thread=2a10001fca0: Panic - Generated at user request
```

Recovery: Reboot the affected domain.

Workaround: Before performing any memory DR operations on a domain that owns an NIU device, configure the domain to plumb all NIU devices at boot time, and then reboot. NIU devices should be plumbed at boot time and remain plumbed before you attempt any memory DR removal operations.

To determine if NIU devices are present on a domain, check the `/etc/path_to_inst` file for references to NIU devices.

```
# grep nxge /etc/path_to_inst
"/niu@80/network@0" 0 "nxge"
"/niu@80/network@1" 1 "nxge"
```

The `/niu@80` string in the two lines indicates that `nxge` instances `0` and `1` are for NIU devices. The `nxge0` and `nxge1` devices should be plumbed at boot time. These devices must remain in the plumbed state when you perform any memory DR operations on the domain by using `ldm rm-mem` or `ldm set-mem` command.

To plumb the `nxge0` and `nxge1` instances without configuring network parameters for the devices, create the `/etc/hostname.nxge0` and `/etc/hostname.nxge1` files as follows:

```
# touch /etc/hostname.nxge0
# touch /etc/hostname.nxge1
```

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

Bug ID 6983279: When multiple NIU `nxge` instances are plumbed on a domain, the `ldm rm-mem` and `ldm set-mem` commands, which are used to remove memory from the domain, might never complete. To determine whether the problem has occurred during a memory removal operation, monitor the progress of the operation with the `ldm list -o status` command. You might have encountered this problem if the progress percentage remains constant for several minutes.

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

Workaround: Cancel the `ldm rm-mem` or `ldm set-mem` command, and check if a sufficient amount of memory was removed. If not, a subsequent memory removal command to remove a smaller amount of memory might complete successfully.

If the problem has occurred on the primary domain, do the following:

1. Start a delayed reconfiguration operation on the primary domain.

```
# ldm start-reconf primary
```
2. Assign the desired amount of memory to the domain.
3. Reboot the primary domain.

If the problem occurred on another domain, stop the domain before adjusting the amount of memory that is assigned to the domain.

Direct I/O Configuration Is Not Properly Rebuilt by the `ldm init-system` Command

Bug ID 6980444: The existing I/O devices on the primary domain are not properly removed when the primary domain is reconfigured by using the `ldm init-system` command with an XML file. This command might cause binding failures on the guest domain. These failures occur if the primary domain is still binding PCIe leaf node devices to the primary domain.

Workaround: Use the `ldm rm-io` command to manually remove the PCIe leaf node devices from the primary domain. Then, reboot the primary domain before you bind and start the guest domains.

Guest Domains Cannot Boot From an Exported DVD Device

Bug ID 6981081: When a bootable physical CD or DVD is exported as a virtual disk, the virtual CD or DVD might not be bootable from the guest domain that uses it. Also, the boot might fail with an error similar to the following:

```
{0} ok boot /virtual-devices@100/channel-devices@200/disk@1:f
Boot device: /virtual-devices@100/channel-devices@200/disk@1:f File and args:
Bad magic number in disk label
ERROR: /virtual-devices@100/channel-devices@200/disk@1: Can't open disk label package
ERROR: boot-read fail
Can't open boot device
```

Whether this problem occurs depends on the type of physical CD or DVD drive that is installed on the system.

Emulex HBA Permanently Destroyed If the Firmware Upgrade Is Interrupted

Bug ID 6950462: The Emulex driver automatically performs a firmware upgrade of the Emulex host bus adapter (HBA) when the driver and HBA firmware levels do not match. This might occur when the Emulex HBA is assigned to an I/O domain by means of the direct I/O feature. If the firmware upgrade of an Emulex HBA is interrupted, the card is permanently destroyed. So, do *not* assign any Emulex HBA to an I/O domain with the direct I/O feature.

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

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

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

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

# ldm add-mem 16G ldom1

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

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

```
# ldm list ldom1
```

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

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

Recovery: Reboot the domain.

Cryptographic Units Cannot Be Removed From the primary Domain

Bug ID 6978843: Sometimes, when you attempt to dynamically remove cryptographic units, the following message is issued:

```
# ldm set-crypto 0 primary
```

```
Aug 20 13:02:27 guest1 ncp: WARNING: ncp0: ncp_mau_unconfig:
unable to find MAU for cpu 112
```

```
Aug 20 13:02:27 guest1 ncp: WARNING: ncp0: ncp_mau_unconfig:
unable to find MAU for cpu 104
```

Workaround: Determine whether any CPUs are faulted, and if they are, mark them as being online.

```
# psrinfo
```

```
# psradm -n 0-127
```

Use delayed reconfiguration to remove the cryptographic units.

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

```
# ldm set-crypto 0 primary
```

```
# reboot
```

SPARC T3-1: Detect And Handle Disks That Are Accessible Through Multiple Direct I/O Paths

Bug ID 6984008: A SPARC T3-1 system can be installed with dual-ported disks, which can be accessed by two different direct I/O devices. In this case, assigning these two direct I/O devices to different domains can cause the disks to be used by both domains and impact each other based on the actual usage of those disks.

Workaround: Do not assign direct I/O devices that have access to the same set of disks to different I/O domains. The steps to determine if you have dual-ported disks on T3-1 system are as follows:

Determine whether the system has dual-ported disks by running the following command on the SP:

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


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

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

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

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

System Powers Off When a Memory Request Results in a Hypervisor Abort

Bug ID 6986076: A hypervisor abort could cause your system to power off if all of the following conditions are met:

- The control domain is in delayed reconfiguration mode.
- The first operation after you enter delayed reconfiguration mode is *not* a memory request. A memory request is an `ldm add-mem`, `ldm set-mem`, or `ldm rm-mem` command.
- You subsequently issue one or more memory requests, and reboot the primary domain.

Workaround: Do one of the following if you plan to issue one or more memory requests to the control domain while it is in delayed reconfiguration mode:

- Ensure that you issue one of the memory requests first.
- If you already issued a non-memory request, cancel the delayed reconfiguration by using the `ldm cancel-reconf primary` command. Then, re-enter delayed reconfiguration mode, and re-issue the memory operation first.

nxge Driver Supports Only Eight Hybrid I/O Shares in a Guest Domain

Bug ID 6805126: Each guest domain that uses hybrid I/O cannot support the maximum number of 32 hybrid I/O shares. If you attempt to configure more than 8 virtual network devices in hybrid mode on a guest domain, the system might panic.

Workaround: Configure no more than eight virtual network devices in hybrid mode on a single guest domain.

Documentation Errata

This section contains documentation errors that have been found too late to resolve for the Oracle VM Server for SPARC 2.0 release.

Incorrect Parameter Names in the Input/Output Bus Table

Bug ID 6843196: “Input/Output Bus Table (IOBusTable)” on page 31 of the *Logical Domains (LDoms) MIB 1.0.1 Administration Guide* shows incorrect parameter names.

IOBusDevName should be IOBusName, and IOBusDevPath should be IOBusPath.

virtinfo(1M) and Related API Man Pages Missing From the Oracle Solaris 10 9/10 Release

The `virtinfo(1M)`, `v12n(3EXT)`, and `libv12n(3LIB)` man pages are missing from the Oracle Solaris 10 9/10 man pages.

For the Oracle VM Server for SPARC 2.0 release, see these man pages in the [Oracle VM Server for SPARC 2.0 Reference Manual](#).

Cannot Upgrade From Logical Domains 1.0 to Oracle VM Server for SPARC 2.0

“Upgrading to Oracle VM Server for SPARC 2.0 Software” in [Oracle VM Server for SPARC 2.0 Administration Guide](#) refers to an upgrade path where you can upgrade an UltraSPARC T1 system that is running the Logical Domains 1.0 software to the Oracle VM Server for SPARC 2.0 software.

You cannot perform this kind of upgrade because the UltraSPARC T1 platform is not supported by the Oracle VM Server for SPARC 2.0 software.

Resolved Issues

This section lists bugs that have been fixed since the previous Oracle VM Server for SPARC (or Logical Domains) software release.

Oracle VM Server for SPARC 2.0 RFEs and Bugs Fixed in Oracle Solaris 10 9/10

The following Logical Domains requests for enhancements (RFEs) and bugs were fixed for the Oracle Solaris 10 9/10 release.

- 6669994 Add a domain service to support OS identification - solaris
- 6680702 RFE: Want a mechanism for Solaris to know if it is an LDOM and if so what type
- 6699505 `vdc_init_attr_negotiation()` should compare `msglen` to the size of `vd_attr_msg_t` not `vio_ver_msg_t`

-
- 6708257 DVD-ROM (not OS installation disk) can not mount from guest domain of LDOM 1.0.2
 - 6712763 RFE: Need information about control domain from within a guest domain
 - 6717615 Ambiguous treatment of offline cpus with respect to DR
 - 6720954 add memory dr feature to ldoms
 - 6726533 vdisk failover should handle storage/storage-path failures (a la mpxio)
 - 6741065 Solaris should support cooperative migration
 - 6794962 cpu DR incorrectly rejects some attempts to remove vcpus
 - 6813200 Logical Domains Agents
 - 6832285 Recursive mutex panic when failures from hv_ldc_rx_set_qhead() are not handled properly.
 - 6840358 i_ldc_rxq_drain should not call i_ldc_set_rx_head if the queue is empty
 - 6856230 service domain panics when exporting a file as a single-slice disk
 - 6873151 RFE: unique ID for LDOMs to facilitate asset tracking
 - 6873667 a domain running Sun Cluster should notify the cluster framework when it is migrating
 - 6875401 In elastic mode, a guest at the reboot menu can stall ldm commands
 - 6915529 User level Domain service loopback client registration fails in some scenarios
 - 6924052 iostat within guest domains report 100% busy following a control domain reboot
 - 6929496 deadlock with mdeg_notify_clients and mdeg_unregister hit with DEBUG, causes login hang
 - 6931722 vdisk using failfast or multipath can hang during heavy I/O
 - 6931950 vdisk does not timeout when backend does not exist
 - 6939707 domain can panic on 'ldm stop'
 - 6940624 Spurious DS console message: Service 'agent-system'/client already registered
 - 6945577 LDoms DS fails to register SP 'pri' if more than 10 active domains booted.

RFEs and Bugs Fixed for Oracle VM Server for SPARC 2.0 Software

The following Oracle VM Server for SPARC 2.0 RFEs and bugs were fixed for the Oracle VM Server for SPARC 2.0 software release:

- 6717661 ldmd does not report MAUs (CPUset) in order.
- 6783343 LDom Manager needs to parse components portion of PRI for rcid
- 6852143 set-vcv prevented when system already in del-reconf mode
- 6887611 ldm ls -o checks are not consistent; some check 2 letters, some check the entire key
- 6903235 RFE: LDoms Manager should support cooperative migration
- 6904672 ldmp2v should support root file system on slices other than s0
- 6905204 “WARNING: ncp1: only one instance (0) allowed” seen on doing P2V migration of a sun4u to N2/VF
- 6908327 ldmp2v prepare usage message could be improved
- 6908985 DRM: an enabled policy may not decrease cpu count when it needs to
- 6909998 DRM: bit of inconsistency in assigning tod-begin and tod-end parameter values
- 6914643 ldmd core dumps on upgrade from 1.2 patch-02 to 1.3 when inactive domains are in existence.
- 6916428 ldmp2v: Flash archiving method should support excluding file systems, directories and files
- 6919211 ldmp2v should be aware of /var/fm
- 6920852 ldmp2v convert can fail if source system has 0.0.0.0 ip address
- 6922211 Unnecessary message in ldmd SMF log: warning: vswitch node missing vlan-id property
- 6923698 /usr/sbin/ldmconfig maximum number of ldoms don't match documentation
- 6924550 Dring full, vdc/vds hang while “ldmp2v prepare” a domain.
- 6926557 RFE: Support new version of mdstore with timestamp and config size
- 6928250 DRM policy handle mismatch with manual
- 6929164 ldmp2v(1M) collect phase should not ufsdump(1M) on actively mounted FS.
- 6932645 PM legacy memory leaks plus clean-up
- 6933260 Mount vdisks on files with -o nologging

-
- 6934296 “ldm migrate-dom” is not correctly interpreted
 - 6936397 DRM: is not reducing vcpu down to 'min-vcpu' on systems that are basically idle (%util=0.1)
 - 6938200 code to get PM policy should not retry every second when SNMP support can't be initialized
 - 6942367 LDoms configuration GUI needs to be updated with Oracle branding
 - 6947814 Domain boot failure if target ldmd dies during a migration
 - 6948564 pmi_get_policy_by_snmp() calls often take over a second to complete
 - 6949238 Guest domain hangs after P2V from Solaris 10 update N to update N
 - 6949242 Migration error for incompatible FW versions is not informative
 - 6950417 when DRM policy times out, vcpu count remains at the same level
 - 6951597 add a resource_id to pcie_device path components
 - 6953092 root disk on EMC Powerpath device confuses ldmp2v
 - 6953460 ldoms seems to be querying HV utilization way to often
 - 6953655 LDoms guest domain MAC address and hostid automatic assignment by ldmd is not very random
 - 6953803 ldmd coredump in ASSERT(snmp_hdlp == NULL)
 - 6957982 CPU Defrag needs to lower its frequency of attempts after repeated failures occur.
 - 6959476 ldmd not properly informing SMF start script to restart the service on certain failures
 - 6961659 rm-vdpcs causes delayed reconfig and later core dumps
 - 6962470 Migration state MD version failures not handled gracefully
 - 6962545 A source ldmd crash during migration causes restart to fail
 - 6966817 ldm(1M) needs an execution profile
 - 6986076 Valid RA's should not change mblocks unless their old mblock is unconfigured
 - 6988211 migration of an inactive domain with a vdsdev causes ldmd to core dump on target machine
 - 6991330 LDoms manager needs to wait longer on MD store requests since GM's host flash driver is slow
 - 6991946 SDIO: ldmd fails to restart due to fatal error on split-pci/sdio config
 - 6995317 SPARC T3-4 server: Migration of active domains fails saving guest state

6996896 prom page locked panic on booting large sparse memory guest domain with LDoms Manager 2.0 patch

6997006 Bound domain memory addition resulted in ldmd SEGV core dump with LDoms Manager 2.0 patch