Oracle® VM Server for SPARC 2.1 Release Notes

Part No: 821–2856–12
February 2012

ORACLE

120217@25097

Contents

Preface

The *Oracle VM Server for SPARC 2.1 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 2.1 software.

Related Documentation

The following table shows the documentation that is available for the Oracle VM Server for SPARC 2.1 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 2.1 Software	*Oracle VM Server for SPARC 2.1 Administration Guide* *Oracle VM Server for SPARC 2.1 Reference Manual* *Oracle VM Server for SPARC 2.1 Release Notes* Oracle Solaris 10 Reference Manual Documentation ■ `drd`(1M) man page ■ `vntsd`(1M) man page	821-2854 821-2855 821-2856
Oracle Solaris OS: Installation and Configuration	Oracle Solaris 10 9/10 Release and Installation Documentation	N/A

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

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

Access to Oracle Support

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

Typographic Conventions

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

TABLE P–2 Typographic Conventions

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

TABLE P–3 Shell Prompts *(Continued)*

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

◆ ◆ ◆ CHAPTER 1

Oracle VM Server for SPARC 2.1 Release Notes

These release notes contain changes for this release, a list of supported platforms, a matrix of required software and patches, and other pertinent information, including bugs that affect the Oracle VM Server for SPARC 2.1 software.

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

Note – Oracle VM Server for SPARC features are added and maintained on the supported hardware platforms listed in "Supported Platforms" on page 10. However, new features will not be added and existing features will not be maintained on hardware platforms that have been removed from the list.

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

What's New in This Release

The major changes for this release of the Oracle VM Server for SPARC 2.1 software are as follows:

- Support for performing named CPU core and memory block assignments (only with the Oracle Solaris 11 Support Repository Update 4 (SRU 4)). See "Assigning Physical Resources to Domains" on page 18.
- Support for Oracle's SPARC T4 servers (only with Oracle Solaris 10 patch ID 147507-01 or Oracle Solaris 11 SRU 4). See "Supported Platforms" on page 10.

- Single-CPU performance enhancements that provide high instructions per cycle (IPC) on the SPARC T4 platform (only with Oracle Solaris 10 patch ID 147507-01 or Oracle Solaris 11 SRU 4). See *Tuning the SPARC CPU to Optimize Workload Performance on SPARC T4 Systems* at Oracle VM Technical White Papers (`http://www.oracle.com/technetwork/server-storage/vm/overview/index.html`).
- Live migration support for active domains. See Chapter 9, "Migrating Domains," in *Oracle VM Server for SPARC 2.1 Administration Guide*.
- Support for a higher-priority domain to take CPUs from a lower-priority domain when required by DRM policy. See the description of the `priority` property in the `ldm`(1M) man page.
- Support for disabling inter-vnet channels. See "Inter-Vnet LDC Channels" in *Oracle VM Server for SPARC 2.1 Administration Guide* and the `ldm`(1M) man page.
- Support for virtual device service validation. See the `ldm`(1M) man page.
- Oracle VM Server for SPARC P2V tool enhancements. See Chapter 13, "Oracle VM Server for SPARC Physical-to-Virtual Conversion Tool," in *Oracle VM Server for SPARC 2.1 Administration Guide*.
- Oracle VM Server for SPARC MIB enhancements. See Chapter 15, "Using the Oracle VM Server for SPARC Management Information Base Software," in *Oracle VM Server for SPARC 2.1 Administration Guide*.
- Extended mapin space support for the Oracle Solaris 10 8/11 OS and Oracle Solaris 11 OS. See the `ldm`(1M) man page.
- Support for the integrated dynamic reconfiguration of cryptographic units and virtual CPUs.
- Bug fixes

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

System Requirements

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

Supported Platforms

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

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

Note – Starting with Logical Domains 1.3, UltraSPARC T1 platforms are no longer supported. Earlier releases of the Logical Domains software continue to support these platforms. By default, the Oracle Solaris 11 OS includes Version 2.1 of the Logical Domains Manager, but not the changes that are provided by the Oracle Solaris 11 SRU 4.

However, Logical Domains Manager 2.1 is *not* supported on UltraSPARC T1 systems.

The Oracle VM Server for SPARC 2.1 software is supported on the following platforms:

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

- Oracle's Sun Blade T6340 server module (refer to the *Sun Blade T6340 Server Module Product Notes*)
- Oracle's Netra T5440 server (refer to the *Sun Netra T5440 Server Product Notes*)
- Oracle's Sun Netra T6340 Server Module (refer to the *Sun Netra T6340 Server Module Product Notes*)

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

Required Software and Patches

This section lists the required software and patches for use with the Oracle VM Server for SPARC 2.1 software.

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

Required and Recommended Oracle Solaris OS

To use all the features of the Oracle VM Server for SPARC 2.1 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.

Note – To use all the features of the Oracle VM Server for SPARC 2.1 software on the SPARC T4 platform, the operating system on all domains should be at least the Oracle Solaris 10 8/11 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 2.1 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		

TABLE 1–1 Patches for Older OS Versions and Domains Needing a Patch *(Continued)*

Patch ID	Control Domain	Service Domain	I/O Domain	Guest Domain
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 2.1 Features

To enable all the Oracle VM Server for SPARC 2.1 features, you must run at least version 7.4.0 of the system firmware on both UltraSPARC T2 servers and UltraSPARC T2 Plus servers. You must run at least version 8.1.0 of the system firmware on SPARC T3 servers. You must run at least version 8.1.*x* of the system firmware on SPARC T4 servers. This firmware is preinstalled on the SPARC T4 servers. For information about the required Oracle Solaris OS, see "Required and Recommended Oracle Solaris OS" on page 12.

Required and Recommended System Firmware Patches

To take advantage of all features of Oracle VM Server for SPARC 2.1, ensure that your server runs at least these revisions of the following system firmware patches:

147307-01 Sun SPARC Enterprise T5120 and T5220 Servers

147308-01 Sun Blade T6320 Server Module

147309-01 Netra T5220 Server

147310-01 Sun SPARC Enterprise T5140 and T5240 Servers

147311-01 Sun SPARC Enterprise T5440 Server

147312-01 Sun Blade T6340 Server Module

147313-01 Netra T5440 Server

147314-01 Sun Netra T6340 Server Module

147315-01 SPARC T3-1 Server

147316-01 SPARC T3-2 Server

147317-01 SPARC T3-4 Server

147318-01 SPARC T3-1B Server

147319-01 Netra SPARC T3-1 Server

147320-01 Netra SPARC T3-1B Server

147284-01	SPARC T4-1 Server
147285-01	SPARC T4-2 Server
147286-01	SPARC T4-4 Server
147287-01	SPARC T4-1B Server
147289-01	Netra SPARC T4-1 Server
147290-01	Netra SPARC T4-2 Server
147292-01	Netra SPARC T4-1BA Server

Note – SPARC T4 systems are preinstalled with the required firmware.

Minimum Version of Software Required

Running the Oracle VM Server for SPARC 2.1 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 2.1 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 2.1 Features" on page 13, and with the Oracle Solaris 10 9/10 OS.

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

- Upgrade the control and service domains to the Oracle Solaris 10 9/10 OS (or to the Oracle Solaris 10 8/11 OS for SPARC T4 systems)
- Apply the patches listed in Table 1–1

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

Following is a matrix of the minimum versions of software required. The Oracle VM Server for SPARC 2.1 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). See the data sheet for your platform at `http://www.oracle.com/technetwork/documentation/sparc-tseries-servers-252697.html`.

TABLE 1–2 Minimum Versions of Software

Supported Server	System Firmware	Oracle Solaris OS
SPARC T4 servers	8.1.*x*	See the data sheet for your SPARC T4 server

TABLE 1–2 Minimum Versions of Software *(Continued)*

Supported Server	System Firmware	Oracle Solaris OS
SPARC T3 servers	8.0.*x*	See the data sheet for your SPARC T3 server
UltraSPARC T2 Plus servers	7.3	See the data sheet for your UltraSPARC T2 Plus server
UltraSPARC T2 servers	7.3	See the data sheet for your UltraSPARC T2 server

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

Direct I/O Hardware and Software Requirements

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

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

  Refer to your platform's hardware documentation to verify which cards can be used on your platform. For an up-to-date list of supported PCIe cards, see Doc 1325454.1: Oracle VM Server for SPARC Direct I/O Feature (`https://support.oracle.com/CSP/main/article?cmd=show&type=NOT&doctype=REFERENCE&id=1325454.1`).
- **Software Requirements.** To use the DIO feature, the following domains must run the supported OS:
  - `primary` **domain.** At least the Oracle Solaris 10 9/10 OS plus patch ID 145868-01
  - **I/O domain.** Any Oracle Solaris OS that is supported by the platform

Live Domain Migration Requirements

The Oracle VM Server for SPARC 2.1 release introduces *live migration*, which provides performance improvements that enable an active domain to be migrated while it continues to run. See Chapter 9, "Migrating Domains," in *Oracle VM Server for SPARC 2.1 Administration Guide*.

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

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

- Version 8.1.0 for SPARC T3 systems
- Version 8.1.*x* for SPARC T4 systems

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

Location of Oracle VM Server for SPARC 2.1 Software

You can find the Oracle VM Server for SPARC 2.1 software to download at `http://www.oracle.com/virtualization/index.html`.

The `OVM_Server_SPARC-2_1.zip` file that you download contains the following:

- Oracle VM Server for SPARC 2.1 software (`SUNWldm.v`)
- The `ldm`(1M), `ldmconfig`(1M), and `ldmd`(1M) man pages in the `SUNWldm.v` package are installed when the package is installed
- Installation script for Oracle VM Server for SPARC 2.1 software (`install-ldm`)
- Oracle VM Server for SPARC Management Information Base (`SUNWldmib`)
- Physical-to-Virtual Conversion Tool (`SUNWldmp2v`)

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

```
OVM_Server_SPARC-2_1/
   Install/
      install-ldm
   Product/
      SUNWldm.v
      SUNWldmib
      SUNWldmp2v
   README.txt
```

Location of Patches

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

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

Location of Documentation

The *Oracle VM Server for SPARC 2.1 Administration Guide*, the *Oracle VM Server for SPARC 2.1 Reference Manual*, and these *Oracle VM Server for SPARC 2.1 Release Notes* can be obtained from:

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

The Oracle VM Server for SPARC 2.1 man pages are installed on your system as part of the `SUNWldm.v` and `SUNWldmp2v` packages.

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

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

Related Software

Optional Software

Oracle VM Server for SPARC Management Information Base (MIB) software can help you enable third-party applications to perform remote monitoring and a few control operations. For more information, see Chapter 15, "Using the Oracle VM Server for SPARC Management Information Base Software," in *Oracle VM Server for SPARC 2.1 Administration Guide*.

Software That Can Be Used With the Logical Domains Manager

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

- **SunVTS** functionality is available in the control domain and guest domains on certain Logical Domains software releases and certain platforms. SunVTS is a validation test suite that provides a comprehensive diagnostic tool that tests and validates Oracle's Sun hardware by verifying the connectivity and proper functioning of most hardware controllers and devices on Oracle's Sun servers. For more information about SunVTS, refer to the *SunVTS 7.0 Software*.
- **Explorer Data Collector** can be used with the Logical Domains Manager software enabled on the control domain. Explorer is a diagnostic data collection tool. The tool comprises shell scripts and a few binary executables. For more information, see the *Oracle Explorer User's Guide*.
- **Oracle Solaris Cluster** software can be used in a guest domain with some restrictions. See the Oracle Solaris Cluster documentation for more information about any restrictions and about the Oracle Solaris Cluster software in general. Starting with the Logical Domains 1.2 release and the Oracle Solaris Cluster 11/09 release, you can also manage logical domains as a resource by using the Oracle Solaris Cluster failover agent.

- **Oracle Enterprise Manager Ops Center** enables you to manage physical and virtual system resources. This solution simplifies resource discovery and monitoring, provides operating system and firmware provisioning, performs comprehensive update and patch management, manages virtual environments such as 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 2.1 software:

- **Sun Integrated Lights Out Manager (ILOM) 3.0** is the system management firmware that you can use to monitor, manage, and configure UltraSPARC T2, UltraSPARC T2 Plus, SPARC T3, and SPARC T4 server platforms. ILOM is preinstalled on these platforms and can be used on supported servers with the Oracle VM Server for SPARC 2.1 software enabled. Refer to the *Sun Integrated Lights Out Manager 3.0 User's Guide* for features and tasks that are common to Oracle's Sun rackmounted servers or blade servers that support ILOM. Other user documents present ILOM features and tasks that are specific to the server platform that you are using. You can find the ILOM platform-specific information within the documentation set that accompanies your system.
- **Netra Data Plane Software Suite** is a complete board software package solution. The software provides an optimized rapid development and runtime environment on top of multithread partitioning firmware for Sun CMT platforms. The Logical Domains Manager contains some `ldm` subcommands (`add-vdpcs`, `rm-vdpcs`, `add-vdpcc`, and `rm-vdpcc`) for use with this software. Refer to the *Netra Data Plane Software Suite 2.0 User's Guide* for more information about this software.

Assigning Physical Resources to Domains

The domain manager automatically selects the physical resources to be assigned to a domain. You can also explicitly choose the physical resources to assign to a domain. Other than for the control domain, this capability is available when the domain runs the Oracle Solaris 10 or Oracle Solaris 11 OS on any platform that supports the Oracle VM Server for SPARC 2.1 software. For the control domain, this capability is *only* available when it runs the Oracle Solaris 11 OS.

Note – This capability is only available with the Oracle Solaris 11 SRU 4.

Resources that you explicitly assign are called *named resources*. Resources that are automatically assigned are called *anonymous resources*.

Caution – The capability to assign named resources is intended to be used only by expert administrators as it requires careful planning to use.

You can explicitly assign physical resources to the control domain and to guest domains. Because the control domain remains active, it might optionally be in delayed reconfiguration mode before you make physical resource assignments. Or, delayed reconfiguration mode is automatically triggered when you make physical assignments. See "Managing Physical Resources on the Control Domain" on page 20. For information about physical resource restrictions, see "Restrictions for Managing Physical Resources on Domains" on page 21.

You can explicitly assign the following physical resources to the control domain and to guest domains:

- **Physical CPUs.** Assign the physical core IDs to the domain by setting the `cid` property. You can set this property by running the following commands:

  ```
  # ldm add-core cid=core-ID[,core-ID[,...]] ldom
  # ldm set-core cid=core-ID[,core-ID[,...]] ldom
  # ldm rm-core [-f] cid=core-ID[,core-ID[,...]] ldom
  ```

  If you specify a core ID as the value of the `cid` property, *core-ID* is explicitly assigned to or removed from the domain.

- **Physical memory.** Assign a set of contiguous physical memory regions to a domain by setting the `mblock` property. Each physical memory region is specified as a physical memory starting address and as a size. You can set this property by running the following commands:

  ```
  # ldm add-mem mblock=PA-start:size[,PA-start:size[,...]] ldom
  # ldm set-mem mblock=PA-start:size[,PA-start:size[,...]] ldom
  # ldm rm-mem mblock=PA-start:size[,PA-start:size[,...]] ldom
  ```

  To assign a memory block to or remove it from a domain, set the `mblock` property. A valid value includes a physical memory starting address (*PA-start*) and a memory block size (*size*) separated by a colon character (`:`).

Note – You *cannot* use dynamic reconfiguration (DR) to move memory or core resources between running domains when you set the `mblock` or `cid` property, respectively. To move resources between domains, ensure that the domains are in a bound or unbound state. For information about managing physical resources on the control domain, see "Managing Physical Resources on the Control Domain" on page 20.

You can use the `ldm list-constraints` command to view the resource constraints for domains. The `physical-bindings` constraint specifies which resource types have been physically assigned to a domain. When a domain is created, the `physical-bindings` constraint is unset until a physical resource is assigned to that domain. By setting the `mblock` property, the `physical-bindings` constraint is set to `memory`. Likewise, by setting the `cid` property, the `physical-bindings` constraint is set to `core`. If both the `cid` and `mblock` properties are set, the `physical-bindings` constraint is set to `core,memory`.

To change the `physical-bindings` constraint for a resource type on the control domain, you *must* first remove all resources of that type by setting the number of resources to zero. Use one of the following methods:

- Set the number of resources to 0 by using the `ldm set-core 0` or `ldm set-mem 0` command.
- Remove all of the specified `physical-bindings` constraints for a particular resource type.

  To remove all named cores and memory blocks, run the `ldm set-core cid=` and `ldm set-mem mblock=` commands, respectively. To remove all anonymous cores and memory blocks, run the `ldm set-core 0` and `ldm set-mem 0` commands, respectively.

  Because the control domain *must* have CPUs and memory allocated, specifying `cid=` or `mblock=` on the control domain returns an error.
- Delete each resource from the domain individually.

If you use the `ldm add-mem` or `ldm set-mem` command to assign multiple physical memory blocks, the addresses and sizes are checked immediately. Also, a domain that has partial cores assigned to it can use the whole-core semantics only if the remaining CPUs of those cores are free and available.

Managing Physical Resources on the Control Domain

Because the control domain is always active, it might optionally be in delayed reconfiguration mode before you make physical resource assignments. When you explicitly assign physical resources, the control domain is automatically placed in delayed reconfiguration mode and the `physical-bindings` constraint is set.

If `physical-bindings=core`, running the `ldm set-core cid=`*core-ID* `primary` command or the `ldm set-vcpu` *CPU-count* `primary` command causes the `physical-bindings` constraint to

be cleared on the next reboot. If the `physical-bindings` constraint is not set to `core`, run the `ldm set-core cid=`*core-ID* `primary` command to set `physical-bindings=core` on the next reboot.

If `physical-bindings=memory`, running the `ldm set-mem` *size* `primary` command causes the `physical-bindings` constraint to be cleared on the next reboot. If the `physical-bindings` constraint is not set to `memory`, run the `ldm set-mem mblock=`*PA-start*`:`*size* `primary` command to set the `physical-bindings` constraint on the next reboot.

Note – When the control domain is in delayed reconfiguration mode, you can perform unlimited memory assignments by using the `ldm add-mem` and `ldm rm-mem` commands on the control domain. However, you can perform only *one* core assignment to the control domain by using the `ldm set-core` command.

Restrictions for Managing Physical Resources on Domains

The following restrictions apply to the assignment of physical resources on domains:

- You *cannot* make physical and non-physical memory bindings or physical and non-physical core bindings in the same domain. However, you *can* have non-physical memory and physical core bindings or non-physical core and physical memory bindings in the same domain.
- When you add a physical resource to a domain, the corresponding resource type becomes constrained as a physical binding.
- Attempts to add individual CPUs to or remove them from a domain where `physical-bindings=core` will fail.
- For unbound resources, the allocation and checking of the resources can *only* occur when you run the `ldm bind` command.
- When removing physical memory from a domain, you must remove the *exact* physical memory block that was previously added.
- Physical memory ranges must *not* overlap.
- You *cannot* use the `ldm add-vcpu -c` or `ldm set-vcpu -c` command to assign a physical resource to a domain.

Upgrading to Oracle VM Server for SPARC 2.1 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 2.1 software, see "Upgrade to the Oracle VM Server for SPARC 2.1 Software" in *Oracle VM Server for SPARC 2.1 Administration Guide*.

Note – Starting with Logical Domains 1.3, UltraSPARC T1 platforms are no longer supported. Earlier releases of the Logical Domains software continue to support these platforms. By default, the Oracle Solaris 11 OS includes Version 2.1 of the Logical Domains Manager, but not the changes that are provided by the Oracle Solaris 11 SRU 4.

However, Logical Domains Manager 2.1 is *not* supported on UltraSPARC T1 systems.

Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 2.1 software.

General Issues

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

I/O MMU Bypass Mode Is No Longer Needed

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

Service Processor and System Controller Are Interchangeable Terms

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

In Certain Conditions, a Guest Domain's 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 `devid` feature of Solaris Volume Manager by adding the following lines to the `/kernel/dr/md.conf` file:**

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

3. **Reboot the guest domain.**

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

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

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

6. **Reboot the guest domain.**

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

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

   These messages are normal and do not report any problems.

Logical Domain Channels and Logical Domains

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

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

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

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

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

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

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

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

- Virtual network
- Virtual disk

- Virtual console

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

15(1) + 9(2) + 8 x 3(3) = 48 LDCs in total

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

15 + 46 + 45 x 11 = 556 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 8/11 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 2.1 Administration Guide*.

Booting a Large Number of Domains

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

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

If unallocated virtual CPUs are available, assign them to the service domain to help process the virtual I/O requests. Allocate 4 to 8 virtual CPUs to the service domain when creating more than 32 domains. In cases where maximum domain configurations have only a single CPU in the service domain, do not put unnecessary stress on the single CPU when configuring and using the domain. The virtual switch (`vsw`) services should be spread over all the network

adapters available in the machine. For example, if booting 128 domains on a Sun SPARC Enterprise T5240 server, create 4 `vsw` services, each serving 32 virtual net (`vnet`) instances. Do not have more than 32 `vnet` instances per `vsw` service because having more than that tied to a single `vsw` could cause hard hangs in the service domain.

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

Memory and swap space usage increases in a guest domain when the `vsw` services used by the domain provides services to many virtual networks (in multiple domains). This is due to the peer-to-peer links between all the `vnet` connected to the `vsw`. The service domain benefits from having extra memory. Four Gbytes is the recommended minimum when running more than 64 domains. Start domains in groups of 10 or less and wait for them to boot before starting the next batch. The same advice applies to installing operating systems on domains. You can reduce the number of links by disabling inter-vnet channels. See "Inter-Vnet LDC Channels" in *Oracle VM Server for SPARC 2.1 Administration Guide*.

Cleanly Shutting Down and Power Cycling a Logical Domains System

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

▼ Power Off a System With Multiple Active Domains

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

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

▼ Power Cycle the System

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

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

Memory Size Requested Might Be Different From Memory Allocated

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

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

Logical Domains Variable Persistence

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

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

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

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

- At the OpenBoot prompt
- Using the Oracle Solaris OS `eeprom(1M)` command
- Using the Logical Domains Manager CLI (`ldm`)
- Modifying, in a limited fashion, from the system controller (SC) using the `bootmode` command, that is, only certain variables, and only when in the `factory-default` configuration

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

In Oracle VM Server for SPARC 2.1 software, there are a few cases where variable updates do not persist as expected:

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

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

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

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

- When running Live Upgrade, perform the following steps:

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

If you modify the time or date on a logical domain, for example using the `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.

ldmp2v convert Command: VxVM Warning Messages During Boot

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

```
Boot device: disk0:a File and args:
SunOS Release 5.10 Version Generic_139555-08 64-bit
Copyright 1983-2009 Sun Microsystems, Inc. All rights reserved.
```

```
Use is subject to license terms.
Hostname: normaal
Configuring devices.
/kernel/drv/sparcv9/vxdmp: undefined symbol ?romp?
WARNING: mod_load: cannot load module ?vxdmp?
WARNING: vxdmp: unable to resolve dependency, module ?misc/ted? not found
/kernel/drv/sparcv9/vxdmp: undefined symbol ?romp?
WARNING: mod_load: cannot load module ?vxdmp?
WARNING: vxdmp: unable to resolve dependency, module ?misc/ted? not found
/kernel/drv/sparcv9/vxio: undefined symbol ?romp?
WARNING: mod_load: cannot load module ?vxio?
WARNING: vxio: unable to resolve dependency, module ?drv/vxdmp? not found
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
WARNING: VxVM vxspec V-5-0-0 vxspec: vxio not loaded. Aborting vxspec load
WARNING: vxspec : CANNOT INITIALIZE vxio DRIVER
NOTICE: VxVM not started
```

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

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

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

Graphical Configuration Assistant Tool Has Been Removed

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

Upgrade Option Not Presented When Using ldmp2v prepare -R

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

```
AVAILABLE DISK SELECTIONS:
       0. c0d0 <SUN-DiskImage-10GB cyl 282 alt 2 hd 96 sec 768>
          /virtual-devices@100/channel-devices@200/disk@0
```

```
       1. c4t2d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
          /pci@400/pci@0/pci@1/scsi@0/sd@2,0
       2. c4t3d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
          /pci@400/pci@0/pci@1/scsi@0/sd@3,0
Specify disk (enter its number)[0]: 0
selecting c0d0
[disk formatted, no defect list found]
format> p


PARTITION MENU:
        0      - change ‘0’ partition
        1      - change ‘1’ partition
        2      - change ‘2’ partition
        3      - change ‘3’ partition
        4      - change ‘4’ partition
        5      - change ‘5’ partition
        6      - change ‘6’ partition
        7      - change ‘7’ partition
        select - select a predefined table
        modify - modify a predefined partition table
        name   - name the current table
        print  - display the current table
        label  - write partition map and label to the disk
        !<cmd> - execute <cmd>, then return
        quit

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

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

partition>
```

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

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

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

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

# ldm add-mem 16G ldom1

# ldm rm-mem 8G ldom1
```

```
Memory removal failed because all of the memory is in use.

# ldm rm-mem 16G ldom1

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

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

Recovery: Reboot the domain.

ldmp2v Command: ufsdump Archiving Method Is No Longer Used

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

Domain Migration Restrictions

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

Version Restrictions for Migration

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

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

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

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

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

  ```
  # ldm migrate ldg1 system2
  The target machine is running an older version of the domain
  manager that does not support the latest migration functionality.
  ```

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

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

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

  ```
  # ldm migrate ldg1 system2
  The source machine is running an older version of the domain manager
  that is not compatible with the version running on the target machine.
  ```

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

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

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

  ```
  # ldm migrate ldg1 system2
  Unable to migrate guest resource state
  Domain Migration of LDom ldg1 failed
  ```

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

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

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

  ```
  # ldm migrate ldg1 system2
  The target machine is running an older version of the System Firmware
  that is not compatible with the version running on the source machine.
  ```

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

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

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

  ```
  # ldm migrate ldg1 system2
  The source machine is running an older version of the System Firmware
  that does not support the latest migration functionality.
  ```

CPU Restrictions for Migration

If the domain to be migrated is running an Oracle Solaris OS version older than the Oracle Solaris 10 9/10 OS, you might see the following message during the migration:

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

The following CPU requirements and restrictions apply:

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

These restrictions also apply when you attempt to migrate a domain that is running in OpenBoot or in the kernel debugger. See "Migrating a Domain That is Running in OpenBoot or in the Kernel Debugger" in *Oracle VM Server for SPARC 2.1 Administration Guide*.

Oracle VM Server for SPARC MIB Issues

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

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

Incorrect ldomCryptoRpReserved Property Value

Bug ID 7042966: The value of the `ldomCryptoRpReserved` property in the cryptographic unit resource pool (`ldomCryptoResourcePool`) erroneously includes the number of cryptographic unit devices that have been assigned to inactive domains.

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

Bug ID 6521530: You receive empty SNMP tables if you query the Oracle VM Server for SPARC MIB 2.1 software using the `snmptable` command with the `-v2c` or `-v3` option. The `snmptable` command with the `-v1` option works as expected.

Workaround: Use the `-CB` option to use only `GETNEXT`, not `GETBULK`, requests to retrieve data. See "Retrieve Oracle VM Server for SPARC MIB Objects" in *Oracle VM Server for SPARC 2.1 Administration Guide*.

Bugs Affecting the Oracle VM Server for SPARC 2.1 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.

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

Bug ID 7117766: The `ldm init-system` command fails to restore the named CPU core constraints for guest domains from a saved XML file.

Workaround: Perform the following steps:

1. Create an XML file for the `primary` domain.

   ```
   # ldm ls-constraints -x primary > primary.xml
   ```

2. Create an XML file for the guest domain or domains.

   ```
   # ldm ls-constraints -x ldom[,ldom][,...] > guest.xml
   ```

3. Power cycle the system and boot a factory default configuration.
4. Apply the XML configuration to the `primary` domain.

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

5. Reboot.
6. Apply the XML configuration to the guest domain or domains.

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

Named Cores Can Power Off All CPUs When in Bind Mode

Bug ID 7111119: You cannot use the `ldm add-core`, `ldm set-core`, and `ldm remove-core` commands when the domain has the elastic policy enabled.

Workaround: Ensure that the domain has the performance policy enabled.

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

Bug ID 7100859: Your system might panic at boot time if you use direct I/O (`ldm remove-io`) to remove multiple PCIe slots from a multi-socket SPARC T-Series system. This occurs when the paths to the PCIe slots are similar to each other (except for the root complex path). The panic might occur after you remove the PCIe slots and then reboot the `primary` domain. For more information about the direct I/O (DIO) feature, see "Assigning PCIe Endpoint Devices" in *Oracle VM Server for SPARC 2.1 Administration Guide*.

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

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

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

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

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

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

Partial Core primary Fails to Permit Whole-Core DR Transitions

Bug ID 7100841: When the `primary` domain shares the lowest physical core (usually `0`) with another domain, attempts to set the whole-core constraint for the `primary` domain fail.

Workaround: Perform the following steps:

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

   ```
   # ldm list -o cpu
   ```

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

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

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

ldmconfig Is Only Supported on Oracle Solaris 10 Systems

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

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

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

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

Bug ID 7071426: A panic might occur during a migration when the domain being migrated has multiple memory blocks that total over 500 Gbytes. Use the `ldm list -o mem` command to determine the amount of memory on the domain.

The panic stack resembles the following:

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

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

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

```
000002a100a5c650 unix:die+9c (30, 2a100a5c930, 6f696e740a232000, 10009, 2a100a5c710, 10000)
000002a100a5c730 unix:trap+75c (2a100a5c930, 0, 0, 10009, 30027b44000, 2a100a5dca0)
000002a100a5c880 unix:ktl0+64 (7022d6dba40, 0, 1, 2, 2, 18a8800)
000002a100a5c9d0 unix:page_trylock+38 (6f696e740a232020, 1, 6f69639927eda164, 7022d6dba40, 13, 1913800)
000002a100a5ca80 unix:page_trylock_cons+c (6f696e740a232020, 1, 1, 5, 7000e697c00, 6f696e740a232020)
000002a100a5cb30 unix:page_get_mnode_freelist+19c (701ee696d00, 12, 1, 0, 19, 3)
000002a100a5cc80 unix:page_get_cachelist+318 (12, 1849fe0, ffffffffffffffff, 3,
0, 1)
000002a100a5cd70 unix:page_create_va+284 (192aec0, 300ddbc6000, 0, 0, 2a100a5cf00, 300ddbc6000)
000002a100a5ce50 unix:segkmem_page_create+84 (18a8400, 2000, 1, 198e0d0, 1000, 11)
000002a100a5cf60 unix:segkmem_xalloc+b0 (30000002d98, 0, 2000, 300ddbc6000, 0, 107e290)
000002a100a5d020 unix:segkmem_alloc_vn+c0 (30000002d98, 2000, 107e000, 198e0d0,
30000000000, 18a8800)
000002a100a5d0e0 genunix:vmem_xalloc+5c8 (30000004000, 2000, 0, 0, 80000, 0)
000002a100a5d260 genunix:vmem_alloc+1d4 (30000004000, 2000, 1, 2000, 30000004020, 1)
000002a100a5d320 genunix:kmem_slab_create+44 (30000056008, 1, 300ddbc4000, 18a6840, 30000056200, 30000004000)
000002a100a5d3f0 genunix:kmem_slab_alloc+30 (30000056008, 1, ffffffffffffffff, 0, 300000560e0, 30000056148)
000002a100a5d4a0 genunix:kmem_cache_alloc+2dc (30000056008, 1, 0, b9, fffffffffffffffe, 2006)
000002a100a5d550 genunix:kmem_cpucache_magazine_alloc+64 (3000245a740, 3000245a008, 7, 6028f283750, 3000245a1d8,
193a880)
000002a100a5d600 genunix:kmem_cache_free+180 (3000245a008, 6028f2901c0, 7, 7, 7, 3000245a740)
000002a100a5d6b0 ldc:vio_destroy_mblks+c0 (6028efe8988, 800, 0, 200, 19de0c0, 0)
000002a100a5d760 ldc:vio_destroy_multipools+30 (6028f1542b0, 2a100a5d8c8, 40, 0, 10, 30000282240)
000002a100a5d810 vnet:vgen_unmap_rx_dring+18 (6028f154040, 0, 6028f1a3cc0, a00,
200, 6028f1abc00)
000002a100a5d8d0 vnet:vgen_process_reset+254 (1, 6028f154048, 6028f154068, 6028f154060, 6028f154050, 6028f154058)
000002a100a5d9b0 genunix:taskq_thread+3b8 (6028ed73908, 6028ed738a0, 18a6840, 6028ed738d2, e4f746ec17d8,
6028ed738d4)
```

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

Removing a Large Number of CPUs From a Guest Domain

Bug ID 7062298: You might see the following error message when you attempt to remove a large number of CPUs from a guest domain:

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

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

CPU Threading Mode Is Not Restored After a Domain Migration Is Canceled

Bug ID 7061265: If you cancel the migration of a domain that has the `threading` property set to `max-ipc`, the `threading` property value is incorrectly restored to `max-throughput` on the domain to be migrated.

Workaround: Manually reset the `threading` property to `max-ipc` on the domain that will be migrated from the source machine.

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

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

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

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

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

Bug ID 7054326: You *cannot* use Solaris hotplug operations to hot remove a PCIe endpoint device after that device is removed from the `primary` domain by using the `ldm rm-io` command. For information about replacing or removing a PCIe endpoint device, see "Making PCIe Hardware Changes" in *Oracle VM Server for SPARC 2.1 Administration Guide*.

install-ldm Hangs When Run By Using an Absolute Path From Another Directory

Bug ID 7050588: If you specify the absolute path to the `install-ldm` command from another directory, the command hangs.

Workaround: Change to the directory in which the `install-ldm` command is installed before you run the command.

```
# cd dirname/OVM_Server_SPARC-2_1/Install
# ./install-ldm
```

ldm add-dev Can Create a Device Alias That is Longer Than Supported by OpenBoot

Bug ID 7044329: If a guest domain has a virtual device with a name that is longer than 31 characters, OpenBoot issues an error message when the domain is started. The device alias that matches the virtual device name is *not* created.

The error message looks similar to the following:

```
Error: device alias name 'mynet12345678901234567890123456789012345678901234567890'
length is greater than 31 chars, device alias not created
```

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

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

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

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

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

  ```
  # ldm add-vdsdev -q /dev/dsk/c3t1d0s2 vol@primary-vds0
  ```

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

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

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

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

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

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

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

The possible values for the `device_validation` property are:

`0` Disable validation for all devices

`1` Enable validation for network devices

`2` Enable validation for disk devices

`3` Enable validation for network and disk devices

`-1` Enable validation for all type of devices, which is the default

When incoming_migration_enabled=false, Outgoing Migrations Fail

Bug ID 7039793: When `incoming_migration_enabled=false` and `outgoing_migration_enabled=true`, outgoing migrations fail with the following message:

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

When `outgoing_migration_enabled=false`, outgoing migrations are expected to fail.

Workaround: Do the following:

1. Set `incoming_migration_enabled=true`.

   ```
   # svccfg -s ldmd setprop ldmd/incoming_migration_enabled=true
   ```

2. Refresh `ldmd`.

   ```
   # svcadm refresh ldmd
   ```

3. Restart `ldmd`.

   ```
   # svcadm restart ldmd
   ```

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

Bug ID 7038650: When a heavily loaded guest domain has a hybrid I/O configuration and you attempt to migrate it, you might see an `nxge` panic.

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

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

Do Not Use the Sun Management Console Software to Monitor an Oracle VM Server for SPARC System

Bug ID 7037495: Using a Sun Management Console to query the CPU status of an Oracle VM Server for SPARC system has the potential to cause data corruption. The corruption is limited to the data structures that the Hypervisor uses to track running domains, and results in the Logical Domains Manager being unable to start. For this reason, do *not* use the Sun Management Console software to monitor Oracle VM Server for SPARC systems.

Workaround: Power cycle the system to use a configuration that is known to be valid.

Incorrect SP Configuration Is Used as the Default

Bug ID 7037295: If the Logical Domains Manager is restarted or the `primary` domain is rebooted after running the `ldm add-spconfig -r` *spconfig* command, the Logical Domains Manager uses the default configuration rather than the specified configuration, *spconfig*. This means that any subsequent configuration modifications are made to the default configuration rather than to the specified configuration, *spconfig*.

Workaround: Set the Logical Domains Manager current configuration by either performing a power cycle or by running the `ldm add-spconfig` *spconfig* command.

All ldm Commands Hang When Migrations Have Missing Shared NFS Resources

Bug ID 7036137: An initiated or ongoing migration, or any `ldm` command, hangs forever. This situation occurs when the domain to be migrated uses a shared file system from another system, and the file system is no longer shared.

Workaround: Make the shared file system accessible again.

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

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

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

1. Initiate a delayed reconfiguration on the control domain.

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

2. Perform any memory reconfiguration operations first.
3. Perform the CPU reconfiguration operations.

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

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

Incorrect Return Status for a Failed CPU DR Operation on a Domain Booted in Single User Mode

Bug ID 7034498: When in single-user mode, attempting to add a virtual CPU to a domain returns a status value of `0`. The status value for this failure should be `1`.

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

Bug ID 7034191: If the system log service, `svc:/system/system-log`, fails to start and does not come online, the Logical Domains agent service will not come online. When the Logical Domains agent service is not online, the `virtinfo`, `ldm add-vsw`, `ldm add-vdsdev`, and `ldm list-io` commands might not behave as expected.

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

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

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

Kernel Deadlock Causes Machine Hang During a Migration

Bug ID 7030045: The migration of an active guest domain might hang and cause the source machine to become unresponsive. When this problem occurs, the following message is written to the console and to the `/var/adm/messages` file:

```
vcc: i_vcc_ldc_fini: cannot close channel 15

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

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

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

Recovery: Perform a powercycle of the source machine.

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

Bug ID 7027105: A `No response` message might appear in the Oracle VM Server for SPARC log when a loaded domain's DRM policy expires after the CPU count has been substantially reduced. The `ldm list` output shows that there are more CPU resources allocated to the domain than is shown in the `psrinfo` output.

Workaround: Use the `ldm set-vcpu` command to reset the number of CPUs on the domain to the value that is shown in the `psrinfo` output.

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

Bug ID 7026160: You perform a domain migration while a DRM policy is in effect. Later, if the DRM policy expires or is removed from the migrated domain, DRM fails to restore the original number of virtual CPUs to the domain.

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

Virtual CPU Timeout Failures During DR

Bug ID 7025445: Running the `ldm set-vcpu 1` command on a guest domain that has over 100 virtual CPUs and some cryptographic units fails to remove the virtual CPUs. The virtual CPUs are not removed because of a DR timeout failure. The cryptographic units are successfully removed.

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

Domain Bind Fails When XML File Has an Invalid Network or Disk Back End

Bug ID 7024499: If you use an XML file to bind a domain with the `ldm bind -i` *xml-file* command, the bind might fail. The failure is due to an invalid network device or disk back-end path even if you use the `-f` or `-q` option. The bind fails when both of the following circumstances are true:

- The XML file references an invalid network device or disk back-end path.
- The service domain that supports the invalid network device or disk back-end path is up and supports `ldmad`.

Although both the `-f` and `-q` options can be specified with the `bind -i` *xml-file* command, these options are ignored.

Workaround: Do the following:

1. Temporarily disable `ldmad` on the service domains that have an invalid device or back end.

   ```
   # svcadm disable ldoms/agents
   ```

2. Re-enable `ldmad` on each service domain where you disabled `ldmad` after the bind.

   ```
   # svcadm enable ldoms/agents
   ```

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

Bug ID 7023216: A domain cannot be migrated if it contains a duplicate MAC address. Typically, when a migration fails for this reason, the failure message shows the duplicate MAC address. However in rare circumstances, this failure message might not report the duplicate MAC address.

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

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

Simultaneous Migration Operations in "Opposite Direction" Might Cause ldm to Hang

Bug ID 7019493: If two `ldm migrate` commands are issued simultaneously in the "opposite direction," the two commands might hang and never complete. For example, an opposite direction situation is one where you simultaneously start a migration on machine A to machine B and a migration on machine B to machine A.

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

Workaround: None.

Removing a Large Number of CPUs From the Control Domain

Bug ID 6994984: Use a delayed reconfiguration rather than dynamic reconfiguration to remove more than 100 CPUs from the `primary` domain. Use the following steps:

1. Use the `ldm start-reconf primary` command to put the control domain in delayed reconfiguration mode.
2. Partition the host system's resources that are owned by the control domain, as necessary.
3. Use the `ldm cancel-reconf` command to undo the operations in Step 2, if necessary, and start over.
4. Reboot the control domain to make the reconfiguration changes take effect.

SPARC T3: Oracle VM Server for SPARC Hangs When Performing Memory Operations

Bug ID 6994300: The Logical Domains Manager might hang on a SPARC T3 system when performing memory operations and possibly migration operations. Such operations will fail to complete.

This hang might occur on any T3 platform that uses any network interface unit (NIU) adapter, but the hang has been confirmed on systems that have XAUI extenders.

Workaround: Apply patch ID 144500-19.

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

Bug IDs 6989192 and 7071760: You might experience OS hangs at login or while executing commands when the following conditions are met:

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

Workaround: Apply patch ID 147149-01.

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

Bug ID 6984681: When using the `pkgadd` command to install the `SUNWldm.v` package from a directory that is exported by means of NFS from a Sun ZFS Storage Appliance, you might see the following error message:

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

Workaround: Ignore this message.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

ldmd Falsely Reports 100% CPU Utilization on a Domain

Bug ID 6982280: In rare instances when in elastic mode, `ldmd` might falsely report that a few CPUs performing I/O on a guest domain are at 100% utilization. This `ldmd` report contradicts the actual processor status that is reported by running `prsinfo` on the guest domain.

Workaround: Set the CPU count on the guest domain to be 2. Then, reset the CPU count to the original value.

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.

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.

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
```

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 machine 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 machine 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
   ```

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

Bug ID 6968507: Following the migration of an active domain, CPU utilization in the migrated domain can increase dramatically for a short period of time. If a dynamic resource managment (DRM) policy is in effect for the domain at the time of the migration, the Logical Domains Manager might begin to add CPUs. In particular, if the `vcpu-max` and `attack` properties were

not specified when the policy was added, the default value of `unlimited` causes all the unbound CPUs in the target machine to be added to the migrated domain.

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

An In-Use MAC Address Can be Reassigned

Bug ID 6968100: Sometimes an in-use MAC address is not detected and it is erroneously reassigned.

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

ldmconfig Cannot Create a Domain Configuration on the SP

Bug ID 6967799: The `ldmconfig` script cannot properly create a stored logical domains configuration on the service processor (SP).

Workaround: Do *not* power cycle the system after the `ldmconfig` script completes and the domain reboots. Instead, perform the following manual steps:

1. Add the configuration to the SP.

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

2. Remove the `primary-with-clients` configuration from the SP.

   ```
   # ldm rm-spconfig primary-with-clients
   ```

3. Power cycle the system.

If you do not perform these steps prior to the system being power cycled, the existence of the `primary-with-client` configuration causes the domains to be inactive. In this case, you must bind each domain manually and then start them by running the `ldm start -a` command. After the guests have booted, repeating this sequence enables the guest domains to be booted automatically after a power cycle.

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

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

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

Memory DR Is Disabled Following a Canceled Migration

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

Dynamic Reconfiguration of MTU Values of Virtual Network Devices Sometimes Fails

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

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

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

Memory DR Is Not Supported With Some Physical Memory Configurations

Bug ID 6912155: In certain supported configurations when all the DIMM slots are not populated in a machine, the resulting physical memory address map is not contiguous and can have address "holes" between successive memory blocks. For such a configuration, memory DR is not supported.

Workaround: To reconfigure memory when memory DR is not supported, do the following:

- `primary` **domain.** Use delayed reconfiguration.
- **Other domains.** Stop the domain, perform the memory reconfiguration, and then restart the domain.

For memory layout information, see your platform's hardware documentation.

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.

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 list-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 Commands Are Slow to Respond When Several Domains Are Booting

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

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

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

Bug ID 6853273: While a system 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

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.

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.

Virtual Network Devices Are Not Created Properly on the Control Domain

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

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

For example:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Explicit Console Group and Port Bindings Are Not Migrated

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

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

Constraint Database Is Not Synchronized to Saved Configuration

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
# svcadm restart vntsd
```

Note – This command will disconnect all active console connections.

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

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

Workaround: Use this command instead:

```
uadmin 2 0
```

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

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

Bug ID 6742805: A domain shutdown or memory scrub can take over 15 minutes with a single CPU and a very large memory configuration. During a shutdown, the CPUs in a domain are used to scrub all the memory owned by the domain. The time taken to complete the scrub can

be quite long if a configuration is imbalanced, for example, a single CPU domain with 512 Gbytes of memory. This prolonged scrub time extends the amount of time it takes to shut down a domain.

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

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

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

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

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.

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 2.1 Administration Guide*.

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

Bug ID 6656033: Simultaneous net installation of multiple guest domains fails on systems that have a common console group.

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

The scadm Command Can Hang Following an SC or SP Reset

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

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

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

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

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

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

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

Workaround: You can ignore these messages.

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

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

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

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

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

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

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

```
primary# svcadm restart fmd
```

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

Bug ID 6540368: This issue is summarized in "Logical Domains Variable Persistence" on page 27 and affects only the control domain.

Cannot Set Security Keys With Logical Domains Running

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

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

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

Behavior of the ldm stop-domain Command Can Be Confusing

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

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

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

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

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.

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.

Documentation Errata

This section contains documentation errors that have been found too late to resolve for the Oracle VM Server for SPARC 2.1 release.

Incorrect Cross Reference to Required Software Information

The section "Software Compatibility" in *Oracle VM Server for SPARC 2.1 Administration Guide* incorrectly refers to information about requirements to obtain the latest features. Instead, refer to "Live Domain Migration Requirements" on page 15.

ldm stop Command Description Is Misleading

The description states that the `ldm stop` command issues a `shutdown` request, while it actually issues a `uadmin()` system call.

To shut down a domain in the most "graceful" manner, perform a `shutdown` or `init` operation in the domain that you want to stop. See the shutdown(1M) or init(1M) man page.

Logical Domains Manager Package Name Incorrect in Upgrade Procedure

The name of the Logical Domains Manager package to install is `SUNWldm.v`. Any `pkgadd` command in the Oracle VM Server for SPARC 2.1 documentation must refer to the `SUNWldm.v` package name.

ILOM load Command Synopsis Uses Incorrect Character

The ILOM `load` command synopsis in "Upgrade System Firmware" in *Oracle VM Server for SPARC 2.1 Administration Guide* incorrectly uses a backslash character (\) to indicate that the entire command must be input on a single line.

When specifying this command, do *not* include the backslash character, and ensure that the entire command is input on a single line.

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 2.1 RFEs and Bugs Fixed in Oracle Solaris 10 8/11 OS

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

6846889	Spurious "ds_ldc_cb: LDC READ event" message seen when rebooting guest domain
6850554	Unplumb hio enabled vnet with linkprop set causes memory leaks
6937993	In non-LDoms env applied 141778-02, "~ + Ctrl-B" & "send brk" (not from ALOM) should not be effective
6938259	vnet and vsw should support RxDringData mode
6941249	Assertion failure in vio_allocb()
6946035	sun4v domain services very slow at boot due to taskq maxalloc issues
6947134	Memory leak in vsw_fdbe_add
6949062	vds prints currently inaccessible (error 30) during control domain boot
6949300	LDC should support more than 64MB of shared memory mapins

6959875	vnet and vsw should setup data cookies correctly in RxDringData mode
6961910	Active guest flag becomes '-t----' if a CPU is added after cpu0 was removed
6965789	Migrated domain hangs if it has CPUs that are offline, faulted or spare
6969013	Problem with solaris/vnetwork
6969263	Core dump on max guest creation while accessing agents/ds framework
6969953	vlds_init_sysevent: can't bind to WARNING on console with max guest config and looped virtinfo calls
6971015	Transmit race condition in RxDringData mode causes channel hang
6971450	virtinfo core dumps on SIGSEGV if libds does not exist
6972633	Migration of Guest with Hybrid I/O enabled vnets panics control domain
6980594	vd_setup_partition_vtoc() leaks flabel
6981081	Guest domains can't boot from exported DVD device on T3 platforms
6982725	vnet panics with NULL pointer in vgen_dringsend
6993072	Deadlock: cycle in blocking chain in ds_snmp
6995232	Panic in vgen_ldcsend_dring() when trying to send VLAN tagged packet
7002294	146018-01 is bad patch
7006394	zpool operation panics LDoms parents: mutex_enter: bad mutex, lp=306b99d4cc0 owner=2a100
7010446	libv12n/v12n_domain_roles returns V12N_ROLE_ROOT for all PCIe endpoint device I/O domain
7011341	vnet driver registers the HV API for hybrid I/O even when hybrid I/O is not configured
7015448	Primary panics when plumbing/configuring vnet ports bound to a tavor-sourced vsw
7020002	Solaris panic on vnet port (re)attach during live migration
7023124	ds_send_msg stuck in infinite loop, causes ldm migrate hang, blocks other ldm commands
7025714	Dispatcher assert fails after domain migration
7029310	Panic in ip_input() during live migration

RFEs and Bugs Fixed for Oracle VM Server for SPARC 2.1 Software

The following Oracle VM Server for SPARC 2.1 RFEs and bugs were fixed for the Oracle VM Server for SPARC 2.1 software release:

6447740	RFE: Logical Domains Manager should validate specified vdsdev & net-dev entries
6517847	SANITY() macro needs to die
6697096	ldmd may dump core if multiple set-vcpu operations are carried out while in delayed reconfiguration mode
6703127	VIO DR operations ignore the -f CLI option.
6797724	VIO DR needs to re-evaluate return values
6806121	ldmd -p accepts negative and > 65535 port numbers
6830730	Starting ldmd in performance mode always causes every resource to be powered up regardless of state
6848114	Enhance the way that ldmconfig discovers available disk space
6855018	ldm ls output columns skewed by util formatting
6856201	Need ldmp2v finish script to perform cleanup at first boot of the guest.
6881811	RFE: Automatically handle crypto units when adding/removing CPUs from a domain
6884970	RFE: DRM needs to be able to re-allocate resources between domains
6895868	RFE: LDOMs MIB lagging behind LDOM Management's new features and/or resource options
6896620	RFE: messages written to /var/svc/log/ldoms-ldmd:default.log should be time stamped
6909527	RFE: Provide an option to add vnet devices without inter-vnet LDC channels
6911013	ldm set-vnet mac-addr=..... could be more verbose for user error
6916405	RFE: LDoms Manager support for Live Migration
6922142	/usr/sbin/ldmconfig -c does not work
6928324	Some document issues in ldm manual and message
6932158	Functions to read/write MDs to/from files should be only compiled in when ZDBG=1

6937420	RFE: add --force support to skip the factory-default configuration check in ldm init-system command
6959243	Improve error msg for migration between machines with identical CPU speeds but different STICK freqs
6961846	RFE: Affinity Stage 2: rm-vcpu DR
6964071	WARNING: Received DS snmp data out of sequence with request
6964708	ldm set-vcons shouldn't core dump
6965182	ldm -V emits incorrect name
6965623	Logical Domains CLI: 'set-domain -i' not working
6966519	A cold migration dry run needs to play nice with the sequencer
6969006	Util event is not generated if cpus are made faulty
6969129	{add,set,remove}-memory operation with non-existing ldom name returned message differs from other cmd
6970038	CPUSET for 'ldm ls -o crypto' is not accurate
6971065	sizeof() should never be used in a call to Xcalloc()
6971630	Remove virtinfo(1M), v12n(3EXT), and libv12n(3LIB) man pages from Logical Domains
6971971	Coding error at ldcs.c/lookup_peer_ldc_endpoint()/target_channel
6972829	init-system CLI with -r to clean up the reboot flow and return a positive ACK from ldmd
6973431	ldm set-vdisk command accepts non-numeric input for timeout=
6974115	ldmp2v prepare -C does not delete ZVOL backends
6974560	RFE: Stage 1 - affinity API
6975322	set-vcc port-range while under delayed-reconf fails to hold new port-range after reboot
6976722	Migration state MD version handling needs improvement
6976766	RFE: Upgrade the Logical Domains MIB to use XML v3
6976958	ldmp2v should accept size modifiers for the -M option
6977065	When re-creating guest from an XML doc, guest should be created using cpu alloc unit and not wcore#
6977136	ldmd SMF service should depend on svc:/network/loopback

6977151	Java base "Configuration Assistant" tool shows incorrect memory units.
6977474	Fix watchdog hack to use proper interface to ldoms data structures
6978606	Inactive/Bound domains allow co-existence of whole-core with enabled DRM resource policy.
6978714	ldmp2v should be resilient against non-searchable directories
6979007	SPARC T3-1: Elastic: guest domain panics on boot: sfmmu_cpu_init()
6979580	ldmconfig fails and exits on serial console without giving explanation
6979870	Logical Domains memory block config is suboptimal
6979922	Remove license from files that are delivered in SVR4 packages
6979942	ldmp2v needs to display correct error message when vntsd is disabled.
6980302	Need to improve ldmp2v collect [-O "<flarcreate options>"] usage.
6980444	SDIO config is not properly rebuilt when 'ldm init-system -r -i filename.xml' is used
6980690	fmd operations hang when a guest domain is in bound state
6981458	list-io command shows UNK state for all PCIe devices always
6982150	‘ldm cancel-operation reconf primary‘ returns failure intermittently
6983596	Unable to add vdisk after ldom 1.3 upgrade from ldom 1.0.3
6983722	‘ldm set-mau 0 primary‘ returns status 0 on failure
6984181	RFE: Add support for DTrace Statically Defined Tracing (SDT) Probes to ldmd
6984902	ldmp2v prepare fails on Solaris > 10
6986076	Valid RA's should not change mblocks unless their old mblock is unconfigured
6987028	DRM should print a warning if the default vcpu count cannot be restored
6988211	Migration of an inactive domain with a vdsdev causes ldmd to core dump on target machine
6988928	ldmd has a few small memory leaks
6989485	Empty slots are populated into guest virtual root complexes
6991330	Logical Domains Manager needs to wait longer on MD store requests since GM's host flash driver is slow
6991579	install-ldm installs SUNWldm while ignoring required SSL/PICL packages leading to breakage
6991586	ldmp2v convert should allow specifying the netboot method

6991946	SDIO: ldmd fails to restart due to fatal error on split-pci/sdio config
6992220	RFE: Reduce bind time by replacing linear search done by md_find_string() with a hashing function
6994716	SDIO: automatically detect and disallow SDIO assignment of cards with bridges
6994910	Same physical disk can be exported as slice and full disk to guest Ldom without any error messages.
6995020	Problem with mgmt/ldom-mgr
6995317	SPARC T3-4: Migration of active domains fails saving guest state
6995693	ldmd segv core dumps on incorrect add-vdisk command
6995748	Guest domain panics on SPARC T3-4 when running PM_07
6996085	ldm {add|set}-policy messages should refer to vcpu-min and vcpu-max (not vcpu_min and vcpu_max)
6996515	Memory DR sanity causes system panic on T5440 & SPARC T3-4 with Logical Domains Manager 2.0 patch.
6996896	Prom page locked panic on booting large sparse memory guest domain with Logical Domains Manager 2.0 patch.
6997006	Bound domain memory addition resulted in ldmd SEGV core dump
6997494	auth_ldom‘priv_off() has its error messages swapped
6998038	Improve pri parsing for pm
6998168	xml_v3_tab has duplicate entries
7000327	RFE: Add LDom Manager support to config extended mapin space
7002565	ldmd SMF service start method checking on /etc/security/prof_attr can lead to maintenance state
7002760	P2V sanity test fails on SPARC T3-2
7003970	XMLv3 interface: UUID is missing from the list-constraints/list-bindings XML responses
7004057	Only in performance mode allow Test Harness commands to set pm state of resources
7004566	Fix package metadata to refer to Virtualization for SPARC T-Series
7005197	Restore shorter mdstore DS timeout on newer minor version of DS
7006071	RFE: Reorganize lint checking to check all errors and only exclude ones which still report warnings

7010058	Some minor corrections needed in the LDOM 2.0 Admin Guide
7011541	ldmd_start should not modify audit files
7012007	Change in low/high water mark time has no effect on frequency of memory utilization events
7012327	Links out of date in the Admin Guide under "using power management"
7012573	Example command under "using power management" in the Admin Guide for LDOM2.0 seems out of date
7013854	When multiple DRM policies expire at the same time they do not all successfully revert
7014211	RFE: option to skip ping test during ldmp2v convert
7014492	'dmake all' files race dirs: some dirs lose some of the time
7015153	RFE: LDom Migration should use multiple connections for increased transfer speed
7015829	Autosave config with autorecovery=notify emits incorrect message in factory-default config.
7016510	Migration blocked if a guest registers too many HV API groups
7016989	ldm ls-devices -a io output is skewed
7018069	Minor cleanup of SDIO constraints processing
7018118	ldmd dumping core on md_find_node_prop when adding or removing devices from SDIO domain
7018662	ldmd coredumping with vcpus beyond capacity plus expected results with maus beyond capacity
7019282	Store DB and initiate delay reconfig in wrong order in set-vsw operation
7019811	ldmd core dump on migration testing
7019842	UMEM_DEBUG=default causes Logical Domains Manager crash when domain is booting/stopping
7022018	Bind of a guest domain failing in HV
7022262	start-reconf CLI should check for ongoing migrations.
7023728	ldmp2v dumps core while auto adjusting file system sizes
7033167	ldmd core dumps when attempting a migration when running on downrev firmware
7043129	ldmd needs an option to disable device validation

RFEs and Bugs Fixed for Oracle VM Server for SPARC 2.1 Software Patch

The following RFEs and bugs were fixed as part of the Oracle VM Server for SPARC 2.1 patch:

6617215	ldmd shows MAUs warning when starting up on system which does not have any MAUs
6756315	domain migration needs to handle multipathed disks correctly
6999051	Whole-core constraint lost after power cycle (no HV MD property)
7011573	Support for low-CPI workloads (aka dynamic threading)
7030027	ldmconfig attempts to configure MAU on T4 platform
7030070	PM does not manage virtual CPUs in Elastic mode after a guest starts/reboots
7030098	MEMDR set-mem 8G primary fails under some memory configurations
7031177	set-vsw second time core dump
7036440	HV mblock is modified at both ends
7043747	'ldm bind' broken in elastic mode after 7011573
7049262	ldmd hit a SEGV and dumped core on continuous CPU DR with high-IPC set on the primary domain
7049323	Changing primary threading to max-ipc in delayed reconfig mode post set-vcpu causes ldmd crash
7049377	Live migration of active domains with max-ipc fail
7050138	ldmd crashes on removing DRM policy from the primary in delayed reconfiguration mode
7050996	Need the lowest numbered physical strand to stay in primary domain as workaround for 7043292
7053829	wcore for primary is getting lost after reboot from delayed reconfig
7056927	When whole-core is broken, threading is ignored and the MD is not saved
7065684	Not all updates in delayed reconfig are being marked
7066447	WARNING: lgrp_minlat_node: malformed MD, no CPUs found in latency group
7069952	cancel-reconf restore the removed vdsdev
7070623	ldmd drops core after "set-vsw inter-vnet-link=off"
7076032	An incorrect change of threading when DR is not enabled can cause an HV abort

7098929	Number of vcpus in the DB is out of sync with LDoms Mgr after max-cores is decreased
7101957	Enabling/disabling whole-core constraint in PM elastic mode causes ldmd to dump core.
7107548	inactive domains can not be bound after applying patch 147507-01/-02
7117253	set-mem reports incorrect minimum in error message.