

Oracle[®] VM Server for SPARC 3.2 Release Notes

ORACLE[®]

Part No: E48735
May 2015

Copyright © 2007, 2015, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support 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.

Copyright © 2007, 2015, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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

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

Accessibilité de la documentation

Pour plus d'informations sur l'engagement d'Oracle pour l'accessibilité à la documentation, visitez le site Web Oracle Accessibility Program, à l'adresse <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Accès au support électronique

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

- Using This Documentation** 7

- 1 Oracle VM Server for SPARC 3.2 Release Notes** 9
 - What's New in This Release 9
 - Oracle VM Server for SPARC 3.2 System Requirements 11
 - Deprecated Oracle VM Server for SPARC Features 11
 - Late-Breaking News 12
 - Known Issues 13
 - General Issues 13
 - Migration Issues 18
 - SR-IOV Issues 29
 - Bugs Affecting the Oracle VM Server for SPARC Software 34
 - Resolved Issues 68

Using This Documentation

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

Product Documentation Library

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

Feedback

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

◆◆◆ CHAPTER 1

Oracle VM Server for SPARC 3.2 Release Notes

These release notes contain changes for this release and other pertinent information, including bugs that affect the Oracle VM Server for SPARC 3.2 software.

For information about the supported hardware and required software and firmware, see [Chapter 1, “Oracle VM Server for SPARC 3.2 System Requirements,”](#) in “Oracle VM Server for SPARC 3.2 Installation Guide”.

Note - Oracle VM Server for SPARC features are added and maintained on the supported hardware platforms listed in “[Supported Platforms](#)” in “[Oracle VM Server for SPARC 3.2 Installation Guide](#)”. However, new features will not be added and existing features will not be maintained on hardware platforms that have been removed from the list.

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

Note - The features that are described in this book can be used with all of the supported system software and hardware platforms that are listed in “[Oracle VM Server for SPARC 3.2 Release Notes](#)”. However, some features are only available on a subset of the supported system software and hardware platforms. For information about these exceptions, see [What's New in This Release](#) and [What's New in Oracle VM Server for SPARC Software](#) (<http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html>).

What's New in This Release

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

The major changes for the Oracle VM Server for SPARC 3.2 software are as follows. Note that each feature is available on all supported platforms except UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 unless otherwise stated. For the list of supported platforms, see [“Supported Platforms”](#) in [“Oracle VM Server for SPARC 3.2 Installation Guide”](#).

- Provide performance counter access control to get, set, and unset a domain's access rights to three groups of performance counters. See [“Oracle VM Server for SPARC 3.2 Administration Guide”](#) and the `ldm(1M)` man page.

This feature is not supported on Fujitsu M10 servers.

- Provide migration security enhancements. See [“Oracle VM Server for SPARC 3.2 Administration Guide”](#), [“Oracle VM Server for SPARC 3.2 Security Guide”](#), and the `ldm(1M)` man page.

This feature is also supported on the UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems.

- Provide live migration performance enhancements, including parallel memory-mapped I/O and optimized firmware copying routines. These enhancements reduce the overall migration time and the time during which the migrated guest domain is suspended. See [Chapter 12, “Migrating Domains,”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#) and [“Migration Requirements for Memory”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

Some of these performance enhancements require at least the Oracle Solaris 11.2 SRU 8 OS and are not supported on the Oracle Solaris 10 OS.

- Update virtual disk management for mpgroups to show and specify the selected path to the virtual disk backend. You can see the path that is currently used and dynamically request a path switch. This feature requires at least the Oracle Solaris 11.2.1.5.0 OS. See [“Configuring Virtual Disk Multipathing”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#) and the `ldm(1M)` man page.

This feature is also supported on the UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems.

- Provide a centralized view of the network devices and network statistics in a physical system or physical domain. Also, provide new link properties for virtual network and switch devices. See [Chapter 11, “Using Virtual Networks,”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#) and the `ldm(1M)` man page.

This feature is also supported on the UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems.

- Provide a large pool of LDC endpoints to assign to domains. See [“Using Logical Domain Channels”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).
- Remove virtual CPU cores and virtual memory from a resource group. See [“Using Resource Groups”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).
- Provide support for private virtual LANs. See [“Using Private VLANs”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

This feature is also supported on the UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems.

- Provide support to dynamically assign PCIe buses. See [Chapter 6, “Creating a Root Domain by Assigning PCIe Buses,”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#).
This feature is supported only on the SPARC M5 and SPARC M6 systems and Fujitsu M10 systems.
- Provide support of FIPS 140-2 mode for migration. See [“FIPS 140-2 Mode for Domain Migration”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#)
This feature is also supported on the UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems.
- Provide support for resilient I/O domains that can continue to run even when their root domains are not available. See [“I/O Domain Resiliency”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#)
- Require that each master domain assigned to a slave domain must have the same failure policy. See [“Configuring Domain Dependencies”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#).
This feature is also supported on the UltraSPARC T2, UltraSPARC T2 Plus, and SPARC T3 systems.
- Bug fixes.

Oracle VM Server for SPARC 3.2 System Requirements

You can find information about the recommended and minimum software component versions to use with the Oracle VM Server for SPARC 3.2 release in [Chapter 1, “Oracle VM Server for SPARC 3.2 System Requirements,”](#) in [“Oracle VM Server for SPARC 3.2 Installation Guide ”](#).

Deprecated Oracle VM Server for SPARC Features

Note - Patch updates are no longer available for software versions older than Oracle VM Server for SPARC 3.0. To receive new and updated software, upgrade to the current Oracle VM Server for SPARC release.

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

- The Oracle Solaris 10 OS is no longer supported for use in an SR-IOV root domain. In addition, SR-IOV I/O domain support in the Oracle Solaris 10 OS is deprecated and will be removed in a future product update.
- The network interface functionality of an Oracle VM Server for SPARC vsw driver was deprecated in Oracle Solaris 11.1. The Oracle VM Server for SPARC vsw driver continues to provide virtual network switching functionality for guest domains. See [“Oracle Solaris 11 Networking Overview”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#).

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

- Setting the threading property to manage single-thread workloads on SPARC T4 platforms is deprecated. By default, domains are created for maximum throughput, and the Oracle Solaris OS automatically uses the critical thread API to optimize for single-thread workloads.
- Using the `-c` option of the `add-vcpu`, `set-vcpu`, and `rm-vcpu` subcommands to manage hard partitioning is deprecated. Instead, use the `add-core`, `set-core`, or `rm-core` subcommand to assign whole cores. Also, use the `add-domain` or `set-domain` subcommand to specify the CPU cap (`max-cores`).
- Using the Hybrid I/O feature is deprecated in favor of using the single-root I/O virtualization (SR-IOV) feature.
- Future Oracle VM Server for SPARC releases will no longer support the running of the Oracle Solaris 10 OS in control domains. You can continue to run the Oracle Solaris 10 OS in guest domains, root domains, and I/O domains when using future releases. Oracle Solaris 10 patches will continue to be provided for Oracle VM Server for SPARC 3.2 based on the documented support policy.
- In future Oracle VM Server for SPARC releases, the `ldmp2v` utility will no longer be updated to address bug fix or enhancement requests. This utility will no longer be supported, but will continue to be included and documented as part of the Oracle VM Server for SPARC software.
- Future Oracle VM Server for SPARC releases will no longer include the `ldmconfig` utility as it runs only on Oracle Solaris 10 control domains. This utility will no longer be updated to address bug fix or enhancement requests.
- Using the `ldm migrate-domain -p filename` command to initiate a non-interactive migration operation is deprecated. Instead, use SSL certificate-based authentication. See [“Configuring SSL Certificates for Migration”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

Late-Breaking News

The following sections include information about features or changes that are not described in the Oracle VM Server for SPARC 3.2 books.

Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 3.2 software.

General Issues

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

Oracle VM Server for SPARC 3.2 `ldmd` Daemon Does Not Start If Multiple Virtual Switches Are Assigned to a Single Network Adapter

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

The Oracle VM Server for SPARC 3.1 software restored the original behavior, which prevents you from assigning multiple virtual switches to a single network adapter. However, if you configured your Oracle VM Server for SPARC 3.0 system to assign multiple virtual switches to a single network adapter, the `ldmd` daemon does not start when you upgrade to Oracle VM Server for SPARC 3.2.

Workaround: Perform the following steps:

1. Temporarily re-enable this capability on your Oracle VM Server for SPARC 3.2 system to enable the `ldmd` daemon to start.

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

2. Update your configuration to assign only one virtual switch to a network device.
3. Disable this capability on your Oracle VM Server for SPARC 3.2 system.

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

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

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

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

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

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

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

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

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

3. **Reboot the guest domain.**
After the domain has booted, the Solaris Volume Manager configuration and metadevices should be available.
4. **Check the Solaris Volume Manager configuration and ensure that it is correct.**
5. **Re-enable the Solaris Volume Manager `devid` feature by removing from the `/kernel/drv/md.conf` file the two lines that you added in Step 2.**
6. **Reboot the guest domain.**

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

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

These messages are normal and do not report any problems.

Oracle Solaris Boot Disk Compatibility

Historically, the Oracle Solaris OS has been installed on a boot disk configured with an SMI VTOC disk label. Starting with the Oracle Solaris 11.1 OS, the OS is installed on a boot disk that is configured with an extensible firmware interface (EFI) GUID partition table (GPT) disk label by default. If the firmware does not support EFI, the disk is configured with an SMI VTOC disk label instead. This situation applies only to SPARC T4 servers that run at least system firmware version 8.4.0, to SPARC T5, SPARC M5, SPARC M6 servers that run at least system firmware version 9.1.0, and to Fujitsu M10 servers that run at least XCP version 2230.

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

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

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

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

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

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

- **Solution 1:** Remove the `gpt` property so that the firmware does not report that it supports EFI.
 1. From the OpenBoot PROM prompt, disable automatic booting and reset the system to be installed.

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

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

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

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

3. Begin the Oracle Solaris 11.1 OS installation.

For example, perform a network installation:

```
ok boot net - install
```

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

1. Write an SMI VTOC label on the disk.

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

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

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

The disk should have no other partitions. For example:

```
format> partition
```

```
partition> print
```

Current partition table (unnamed):

Total disk cylinders available: 14087 + 2 (reserved cylinders)

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

3. Re-write the SMI VTOC disk label.

```
partition> label
```

- ```

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

```
- Configure your Oracle Solaris Automatic Installer (AI) to install the Oracle Solaris OS on slice 0 of the boot disk.  
Change the <disk> excerpt in the AI manifest as follows:

```

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

```
  - Perform the installation of the Oracle Solaris 11.1 OS.

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

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

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

```

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

primary# ldm add-mem 16G ldom1

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

primary# ldm rm-mem 16G ldom1

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

```

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

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

- Stop the domain, remove the memory, and then restart the domain.

- Reboot the domain, which causes the Oracle Solaris OS to reallocate its memory management metadata such that the previously added memory can now be removed dynamically in smaller chunks.

## Migration Issues

### Live Migration Might Result In Memory Corruption or Lost Kernel Panic Crash Dumps

**Bug ID 20612716:** A live migration of a guest domain that runs Oracle Solaris 11.2 SRU 8 from a machine with firmware based on Hypervisor 1.14.x to a machine with Hypervisor 1.13.2 might result in memory corruption or lost kernel panic crash dumps after the guest is rebooted.

This problem affects the following live migrations:

- For SPARC T4-based systems, this failure occurs when migrating from a system that runs firmware version 8.7.x to a system that runs up to firmware version 8.6.x.
- For SPARC T5-based systems and other systems that use the 9.x firmware, this failure occurs when migration from a system that runs firmware version 9.4.x to a system that runs up to firmware version 9.3.x.

---

**Note** - Because of the related bug 20594568, you should use this workaround when performing a live migration from any system that has firmware with Hypervisor 1.14.x to any system that has firmware with Hypervisor 1.13.x:

- From a system that runs firmware version 8.7.x to a system that runs up to firmware version 8.6.x
- From a system that runs firmware version 9.4.x to a system that runs up to firmware version 9.3.x

---

**Workaround:** To avoid the problem, add the following line to the `/etc/system` file on the domain being migrated:

```
set retained_mem_already_checked=1
```

For information about correctly creating or updating `/etc/system` property values, see [“Updating Property Values in the /etc/system File”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

Then, reboot the domain before attempting to migrate from Hypervisor version 1.14.x to Hypervisor version 1.13.2.

If the guest domain has been migrated already from firmware 8.7.x to 8.6.x or from 9.4.x to 9.3.x, stop and restart the guest domain. For example:

```
primary# ldm stop-domain domainname
primary# ldm start-domain domainname
```

## A Live Migration of an Oracle Solaris 11.2 SRU 8 Guest Domain to a Target Machine With Version 1.13.1 Hypervisor Is Blocked

**Bug ID 20594568:** A live migration of a guest domain that runs Oracle Solaris 11.2 SRU 8 from a machine with firmware based on Hypervisor 1.14.x to a machine with Hypervisor 1.13.1 is blocked.

```
primary# ldm migrate ldg0 target-host
Target Password:
API group 0x11d v1.0 is not supported in the version of the firmware
running on the target machine.
Domain ldg0 is using features of the system firmware that
are not supported in the version of the firmware running on
the target machine.
```

This problem affects the following live migrations:

- For SPARC T4-based systems, this failure occurs when migrating from a system that runs firmware version 8.7.x to a system that runs up to firmware version 8.5.x.
- For SPARC T5-based systems and other systems that use the 9.x firmware, this failure occurs when migration from a system that runs firmware version 9.4.x to a system that runs up to firmware version 9.2.1.c.

---

**Note** - Because of the related bug 20612716, you should use this workaround when performing a live migration from any system that has firmware with Hypervisor 1.14.x to any system that has firmware with Hypervisor 1.13.x:

- From a system that runs firmware version 8.7.x to a system that runs up to firmware version 8.6.x
- From as system that runs firmware version 9.4.x to a system that runs up to firmware version 9.3.x

---

**Workaround:** To avoid the problem, add the following line to the `/etc/system` file on the domain being migrated:

```
set retained_mem_already_checked=1
```

For information about correctly creating or updating `/etc/system` property values, see [“Updating Property Values in the /etc/system File”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

Then, reboot the domain and retry the migration.

## Domain Migration Might Fail Even Though Sufficient Memory in a Valid Layout Is Available on the Target System

**Bug ID 20453206:** A migration operation might fail even if sufficient memory in a valid layout is available on the target system. Memory DR operations might make it more difficult to migrate a guest domain.

**Workaround:** None.

## Cannot Perform a Live Migration of a Guest Domain That Uses iSCSI Devices

**Bug IDs 19163498 and 16585085:** A logical domain that uses iSCSI devices cannot use live migration.

## Kernel Zones Block Live Migration of Guest Domains

**Bug ID 18289196:** On a SPARC system, a running kernel zone within an Oracle VM Server for SPARC domain will block live migration of the guest domain if it runs one or more “down-revision” components. The following error message is displayed:

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

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

- Stop running the kernel zone.  

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

```
zoneadm -z zonename suspend
```

## Oracle Solaris 10: Domains That Have Only One Virtual CPU Assigned Might Panic During a Live Migration

**Bug ID 17285751:** On the Oracle Solaris 10 OS, migrating a domain that has only one virtual CPU assigned to it might cause a panic on the guest domain in the function `pg_cmt_cpu_fini()`.

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

## Virtual Network Hang Prevents a Domain Migration

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

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

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

---

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

---

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

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

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

**Workaround:** Update the system firmware on the SPARC T4 system to version 8.4. See the workaround for “[Guest Domain Panics at `lgrp\_lineage\_add\(mutex\_enter: bad mutex, lp=10351178\)`”](#) on page 42.

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

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

**Workaround:** None.

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

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

**Recovery:** Restart the logical domain.

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

## `ldm list -o status` on Target Control Domain Reports Bogus Migration Progress

**Bug ID 15819714:** In rare circumstances, the `ldm list -o status` command reports a bogus completion percentage when used to observe the status of a migration on a control domain.

This problem has no impact on the domain that is being migrated or on the `ldmd` daemons on the source or target control domains.

**Workaround:** Run the `ldm list -o status` command on the other control domain that is involved in the migration to observe the progress.

## Oracle Solaris 10: Primary or Guest Domain Panics When Unbinding or Migrating a Guest Domain That Has Hybrid I/O Network Devices

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

**Recovery:** Restart the affected domain.

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

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

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

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

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

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

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

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

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

This issue only applies to UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 systems.

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

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

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

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

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

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

000002a100a5c650 unix:die+9c (30, 2a100a5c930, 6f696e740a232000, 10009,
2a100a5c710, 10000)
000002a100a5c730 unix:trap+75c (2a100a5c930, 0, 0, 10009, 30027b44000,
2a100a5dca0)
000002a100a5c880 unix:ktl0+64 (7022d6dba40, 0, 1, 2, 2, 18a8800)
000002a100a5c9d0 unix:page_trylock+38 (6f696e740a232020, 1, 6f69639927eda164,
7022d6dba40, 13, 1913800)
000002a100a5ca80 unix:page_trylock_cons+c (6f696e740a232020, 1, 1, 5,
7000e697c00, 6f696e740a232020)
000002a100a5cb30 unix:page_get_mnode_freelist+19c (701ee696d00, 12, 1, 0, 19, 3)
000002a100a5cc80 unix:page_get_cachelist+318 (12, 1849fe0, ffffffff, 3,
0, 1)
000002a100a5cd70 unix:page_create_va+284 (192aec0, 300ddbc6000, 0, 0,
2a100a5cf00, 300ddbc6000)
000002a100a5ce50 unix:segkmem_page_create+84 (18a8400, 2000, 1, 198e0d0, 1000,
11)
000002a100a5cf60 unix:segkmem_xalloc+b0 (30000002d98, 0, 2000, 300ddbc6000, 0,
107e290)
000002a100a5d020 unix:segkmem_alloc_vn+c0 (30000002d98, 2000, 107e000, 198e0d0,
30000000000, 18a8800)
000002a100a5d0e0 genunix:vmem_xalloc+5c8 (30000004000, 2000, 0, 0, 80000, 0)
000002a100a5d260 genunix:vmem_alloc+1d4 (30000004000, 2000, 1, 2000,
30000004020, 1)
000002a100a5d320 genunix:kmem_slab_create+44 (30000056008, 1, 300ddbc4000,
18a6840, 30000056200, 30000004000)
000002a100a5d3f0 genunix:kmem_slab_alloc+30 (30000056008, 1, ffffffff,
0, 300000560e0, 30000056148)
000002a100a5d4a0 genunix:kmem_cache_alloc+2dc (30000056008, 1, 0, b9,
fffffffffffffe, 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.

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

**Bug ID 15710957:** When a heavily loaded guest domain has a hybrid I/O configuration and you attempt to migrate it, you might see an nxge panic.

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

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

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

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

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

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

**Bug ID 15701865:** If you attempt a live migration of a domain that depends on an inactive domain on the target machine, the `ldmd` daemon faults with a segmentation fault, and the domain on the target machine restarts. Although you can still perform a migration, it will not be a live migration.

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

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

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

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

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

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

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

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

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

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

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

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

**Workaround:** None.

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

**Bug ID 15655513:** Following the migration of an active domain, CPU utilization in the migrated domain can increase dramatically for a short period of time. If a dynamic resource management (DRM) policy is in effect for the domain at the time of the migration, the Logical Domains Manager might begin to add CPUs. In particular, if the `vcpu-max` and `attack` properties were not specified when the policy was added, the default value of `unlimited` causes all the unbound CPUs in the target machine to be added to the migrated domain.

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

## Memory DR Is Disabled Following a Canceled Migration

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

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

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

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

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

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

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

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

- Install the latest Oracle VM Server for SPARC software on the target machine.

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

## Explicit Console Group and Port Bindings Are Not Migrated

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

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

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

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

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

**Workaround:** When you are configuring the target domain to receive a migrated domain, ensure that the disk volume (vdsdev) matches the disk back end used on the source domain.

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

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

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

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

This failure occurs because migrating the specific memory ranges in use by the source domain requires that compatible memory ranges are available on the target as well. When no such compatible memory range is found for any memory range in the source, the migration cannot proceed. See [“Migration Requirements for Memory”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

**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.

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

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

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

```
svcadm restart vntsd
```

---

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

---

## SR-IOV Issues

### `ldm remove-io` Command Reports a Timeout and Fails to Remove the Last SR-IOV Virtual Function From an I/O Domain

**Bug ID 20731016:** When you use the `ldm remove-io` command to remove the last SR-IOV virtual function from an I/O domain, the command might report a timeout and fail to remove the virtual function.

**Workaround:** If this problem occurs, perform the following steps:

1. Verify that the `system/management/hwmgmt` package is installed on the system.

```
pkg info system/management/hwmgmt
```

2. Disable the `svc:/system/sp/management` service.

```
svcadm disable -st svc:/system/sp/management
```

3. Retry the `ldm remove-io` command.

4. When the SR-IOV virtual function is successfully removed, enable the `svc:/system/sp/management` service.

```
svcadm enable svc:/system/sp/management
```

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

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

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

**Workaround:** None.

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

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

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

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

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

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

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

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

```
set px:pxtool_cfg_delay_usec=25000
```

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

```
svcadm refresh picl
```

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

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

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

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

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

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

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

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

```
VF3: PF has failed
```

```
Mcxnex: HW2SW_MPT command @ failed: 0000ffff
```

```
Hermon: MAD_IFC (port 01) command failed: 0000ffff
```

```
WARNING: mcxnex0: Device Error: HCR Timeout waiting for command go bit
```

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

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

- Perform a reboot of the domain

```
primary# ldm stop -r domain-name
```

- Perform a reset in the SP

```
-> reset /SYS
```

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

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

```
WARNING: kmem_cache_destroy: 'px0_emlxs3_3_cache2'
```

```
(3000383e030) not empty
```

```
WARNING: vmem_destroy('px0_emlxs3_3_vmem_top'):
```

```
leaked 262144 identifiers
```

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

**Workaround:** None.

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

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

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

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

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

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

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

On a Fujitsu M10 server you can assign PCIe endpoint devices and SR-IOV virtual functions from a particular PCIe bus to a maximum of 24 domains. The maximum is 15 domains for supported SPARC T-Series and SPARC M-Series platforms.

## InfiniBand SR-IOV Issues



---

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

---

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

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

## Misleading Messages Shown For InfiniBand SR-IOV Operations

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

Dynamic SR-IOV is not supported for InfiniBand devices.

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

- [“How to Create an InfiniBand Virtual Function”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#)
- [“How to Destroy an InfiniBand Virtual Function”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#)
- [“How to Remove an InfiniBand Virtual Function From an I/O Domain”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#)
- [“How to Remove an InfiniBand Virtual Function From a Root Domain”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide ”](#)

## Bugs Affecting the Oracle VM Server for SPARC Software

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

---

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

---

## Bugs Affecting the Oracle VM Server for SPARC 3.2 Software

### **rcm\_daemon Might Emit a Message on the Console During an `ldm remove-io` Operation**

**Bug ID 20619894:** If the `system/management/hwmgmtd` package is not installed, a dynamic bus remove operation causes the `rcm_daemon` to print the following message on the console:

```
rcm_daemon[839]: rcm script ORCL,pcie_rc_rcm.pl: svcs: Pattern 'sp/management'
doesn't match any instances
```

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

## Recovery Mode Might Hang If the Power Management Policy Is Set to `elastic`

**Bug ID 20570207:** When the power management policy is set to `elastic`, the primary domain might hang while the Logical Domains Manager is recovering domains after detecting faulty or missing resources.

**Recovery:** Change the policy to `disabled` and then power cycle the system to restart recovery mode.

## Unexpected Behavior When Fujitsu M10 Servers Socket Commands Are Used While in Delayed Reconfiguration

**Bug ID 20432421:** If you use the `grow-socket` or `shrink-socket` command to modify virtual CPUs or cores during a delayed reconfiguration, you might experience unexpected behavior. Memory that belongs to the primary domain is reassigned so that only memory in the specified socket is bound to the domain.

**Workaround:** Only modify virtual CPUs or cores by using the `shrink-socket` and `grow-socket` commands while not in a delayed reconfiguration.

## After Dropping Into `factory-default`, Recovery Mode Fails If the System Boots From a Different Device Than the One Booted in the Previously Active Configuration

**Bug ID 20425271:** While triggering a recovery after dropping into `factory-default`, recovery mode fails if the system boots from a different device than the one booted in the previously active configuration. This failure might occur if the active configuration uses a boot device other than the `factory-default` boot device.

**Workaround:** Perform the following steps any time you want to save a new configuration to the SP:

1. Determine the full PCI path to the boot device for the primary domain.  
Use this path for the `ldm set-var` command in Step 4.
2. Remove any currently set `boot-device` property from the primary domain.  
Performing this step is necessary only if the `boot-device` property has a value set. If the property does not have a value set, an attempt to remove the `boot-device` property results in the `boot-device not found` message.

```
primary# ldm rm-var boot-device primary
```

3. Save the current configuration to the SP.

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

4. Explicitly set the boot-device property for the primary domain.

```
primary# ldm set-var boot-device=value primary
```

If you set the boot-device property after saving the configuration to the SP as described, the specified boot device is booted when recovery mode is triggered.

**Recovery:** If recovery mode has already failed as described, perform the following steps:

1. Explicitly set the boot device to the one used in the last running configuration.

```
primary# ldm set-var boot-device=value primary
```

2. Reboot the primary domain.

```
primary# reboot
```

The reboot enables the recovery to proceed.

### **ldm list-rsrc-group Incorrectly Shows PCIe Bus Information Under the Wrong Resource Group**

**Bug ID 20426593:** The `ldm list-rsrc-group` might show I/O resource information under the incorrect resource group when the numerical suffix of the resource group name has more than one digit.

In the following example, the `ldm list-rsrc-group` command incorrectly shows the PCIe bus information for `/SYS/CMIOU10` under the `/SYS/CMIOU1` resource group.

```
primary# ldm list-io
NAME TYPE BUS DOMAIN STATUS
---- -
..
/SYS/CMIOU10/PCIE2 PCIE pci_50 primary OCC
/SYS/CMIOU10/PCIE3 PCIE pci_51 primary OCC
/SYS/CMIOU10/PCIE1 PCIE pci_53 primary OCC
..
.
primary# ldm list-rsrc-group -l -o io /SYS/CMIOU1
NAME
/SYS/CMIOU1

IO
DEVICE PSEUDONYM BOUND
pci@305 pci_5 alt-root
pci@306 pci_6 primary
pci@308 pci_8 alt-root
pci@309 pci_9 primary
```

```
pci@332 pci_50 primary
pci@333 pci_51 primary
pci@335 pci_53 primary
```

PCIe busses pci\_50, pci\_51, and pci\_53 are incorrectly shown under the /SYS/CMI0U1 resource group instead of the /SYS/CMI0U10 resource group.

**Workaround:** Run the `ldm list-io -l` command to obtain the correct resource group name for the PCIe bus from the I/O name. For example, the PCIe bus with the I/O name /SYS/CMI0U10/PCIE2 should belong to /SYS/CMI0U10 and not /SYS/CMI0U1.

## Recovery Mode Does Not Recover a Guest Domain When the Virtual Disk Back End Is Unavailable in the Service Domain

**Bug ID 20321459:** If a virtual disk back end is missing and cannot be validated, the Logical Domains Manager does not recover a guest domain that is assigned the back end. This is true even if multipathing is configured.

**Workaround:** Perform the following steps:

1. Temporarily disable device validation.

```
primary# svccfg -s ldmd setprop ldmd/device_validation integer: 0
primary# svcadm refresh ldmd
primary# svcadm restart ldmd
```

2. Recover the guest domains that are missing the back end manually.

Note that when device validation is disabled, the Logical Domains Manager adds a virtual device to a guest domain even if the back end or associated physical network device does not exist. Thus, ensure that you re-enable device validation after you have recovered the domain configuration.

```
primary# svccfg -s ldmd setprop ldmd/device_validation integer: -1
primary# svcadm refresh ldmd
primary# svcadm restart ldmd
```

## ldm bind Command Issues an Invalid response Error

**Bug ID 20307560:** If you create a guest domain that uses any number of virtual CPUs and any amount of memory and run the `ldm bind` command, the command might issue an Invalid response error. This error might occur if the primary domain has all of the resources before you create the guest domain and run the `ldm bind` command.

**Workaround:** Remove some memory from the primary domain and then run the `ldm bind` command.

## Physical Function or Virtual Function Stuck in the INV State When in Delayed Reconfiguration and Logical Domains Manager Is Restarted

**Bug ID 20257979:** One of the methods to create virtual functions from a physical function is to place the root domain that owns the physical function into delayed reconfiguration. When in the delayed reconfiguration, you can create one or more virtual functions by using the `ldm create-vf` command.

Normally, an `ldm list-io` command shows that the physical function and its child virtual functions are in a clean state. However, if the `ldmd` service is restarted before the root domain is rebooted, or if the delayed reconfiguration is cancelled, the physical function and its virtual functions are marked with the INV state.

The same problem occurs when virtual functions are being destroyed while in delayed reconfiguration. When destroying virtual functions, restarting the Logical Domains Manager and then running the `ldm list-io` output shows no physical functions for that root domain.

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

- Cancel the delayed reconfiguration.  
When you next run the `ldm list-io` command, the physical function and any of its existing virtual functions are in a valid state.
- Reboot the root domain that was in delayed reconfiguration.  
Note that any modifications that you performed while the root domain was in delayed reconfiguration will be present in the OS on the guest domain.

## CPUs Are Not Resumed When the Power Cap Is Under the Limit

**Bug ID 20187197:** If power capping is enabled, sometimes the lowest power state cannot be set. The power state was lowered, but not to the lowest state. When this occurs, the highest power state might not be resumed after setting a higher power limit that warrants the highest power state.

This situation occurs when setting a new power cap limit that is close to the minimum power limit for the system or when setting a new power cap limit where the difference between the actual power (when not power capped) and the new limit would cause the lowest power state to be used.

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

- Disable the power cap
- Set a new power cap limit that is not large or close to the minimum power limit for the system

### **ixgbev Nodes on an I/O Domain Are Reported as Disabled By the ipadm Command And Non-Existent By the ifconfig Command**

**Bug ID 20004281:** When a primary domain is power cycled, ixgbev nodes on the I/O domain might be reported as disabled by the ipadm command, and as nonexistent by the ifconfig command.

**Workaround:** Re-enable the IP interfaces:

```
svcadm restart network/physical:default
```

### **HGXE Interfaces Are Unusable If Assigned by Using Direct I/O to an I/O Domain**

**Bug ID 19943809:** The hxge driver cannot use interfaces inside an I/O domain when the card is assigned by using the direct I/O feature.

The following warning is issued to the system log file:

```
WARNING: hxge0 : <== hxge_setup_mutexes: failed 0x1
```

**Workaround:** Add the following line to the /etc/system and reboot:

```
set px:px_force_intx_support=1
```

### **Guest Domain eeprom Updates Are Lost if an ldm add-spconfig Operation Is Not Complete**

**Bug ID 19932842:** An attempt to set an OBP variable from a guest domain might fail if you use the eeprom or the OBP command before one of the following commands is completed:

- ldm add-spconfig
- ldm remove-spconfig
- ldm set-spconfig
- ldm bind

This problem might occur when these commands take more than 15 seconds to complete.

```
/usr/sbin/eeprom boot-file\=-k
promif_ldom_setprop: promif_ldom_setprop: ds response timeout
eeprom: OPROMSETOPT: Invalid argument
boot-file: invalid property
```

**Recovery:** Retry the eeprom or OBP command after the ldm operation has completed.

**Workaround:** Retry the eeprom or OBP command on the affected guest domain. You might be able to avoid the problem by using the ldm set-var command on the primary domain.

## Rebooting a Guest Domain With More Than 1000 Virtual Network Devices Results in a Panic

**Bug ID 19449221:** A domain can have no more than 999 virtual network devices (vnets).

**Workaround:** Limit the number of vnets on a domain to 999.

## Oracle VM Server for SPARC No Longer Keeps Track of Freed MAC Addresses

**Bug ID 19078763:** Oracle VM Server for SPARC no longer keeps track of freed MAC addresses. MAC addresses are now allocated by randomly selecting an address and then confirming that address is not used by any logical domains on the local network.

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

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

**Workaround:** None.

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

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

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

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

**Workaround:** None.

## Issues Might Arise When FMA Detects Faulty Memory

**Bug ID 17576087:** Performing a power cycle of the system to a saved configuration might not restore the memory after the faulty memory has been replaced.

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

### **DLMP Does Not Work in a Guest Domain on Top of a Virtual Network Device or an SR-IOV Virtual Function**

You cannot configure a DLMP aggregation on an SR-IOV NIC virtual function or a virtual network device in a guest domain.

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

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

```
cannot label 'c1d0': try using fdisk(1M) and then provide a specific slice
Unable to build pool from specified devices: invalid vdev configuration
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Perform the following steps on each domain:

1. Access the console of the domain.

```
primary# telnet localhost 5000
```

2. Set the `boot-device` property.

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

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

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

```
ok> printenv
```

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

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

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

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

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

For information about how to remove PCIe slots from a non-primary root domain, see [“Non-primary Root Domains Overview”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

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

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

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

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

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

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

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

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

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

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

- Ensure that when you assign PCI devices to a guest domain, the card is located physically in the same DCU as the cores.
- Manually assign cores for physical card placement flexibility.

As an example, for a PCI device on IOU0 (pci\_0 through pci\_15), choose a core between 0 and 127, and allocate it to the domain.

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

View the system cores by using the following command:

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

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

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

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

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

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

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

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

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

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

Specifically, the limit may need adjustment if the system is partitioned into multiple logical domains and if too many I/O devices are assigned to any guest domain. Oracle VM Server for SPARC divides the total interrupts into smaller sets given to guest domains. If too many I/O devices are assigned to a guest domain, its supply might be too small to give each device the default limit of interrupts. Thus, it exhausts its supply before it completely attaches all the drivers.

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

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

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

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

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

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

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

The `::irm_pools` and `::irm_reqs` macros show the following information:

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

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

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

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

```
set ddi_msix_alloc_limit = 6
```

When the system successfully boots without warnings, the ::irmpools and ::irmreqs macros show the following updated information:

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

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

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

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

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

This panic occurs if the following conditions exist:

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

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

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

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

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

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

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

2. Shutdown the domain.
3. Stop the host.

```
-> stop /HOST
```

4. Start the host.

```
-> start /HOST
```

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

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

**Workaround:** None.

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

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

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

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

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

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

```
000002a1075c3630 ldc:ldc_mem_rdwr_cookie+20 (c4005fa01c80,
c4004e2c2000,2a1075c37c8, 6c80000, 1, 0)
%l0-3: 00000000001356a4 0000000000136800 0000000000000380
00000000000002ff
%l4-7: 00000000001ad3f8 0000000000000004 00000000ffbf9c
0000c4005fa01c88
000002a1075c3710 vldc:i_vldc_ioctl_write_cookie+a4 (c4004c400030,
380,ffbf98, 100003, 0, 70233400)
%l0-3: 00000000006c80000 0000000000156dc8 0000000000000380
0000000000100003
%l4-7: 00000000702337b0 000002a1075c37c8 0000000000040000
0000000000000000
000002a1075c37f0 vldc:vldc_ioctl+1a4 (3101, c4004c400030,
ffbf98,c4004c400000, c4004c438030, 0)
%l0-3: 0000000000100003 0000000000000000 000000007b340400
0000c4004c438030
%l4-7: 0000c4004c400030 0000000000000000 0000000000000000
0000000000000000
000002a1075c38a0 genunix:fop_ioctl+d0 (c4004d327800, 0, fbf98,
100003,c4004384f718, 2a1075c3acc)
%l0-3: 0000000000003103 0000000000100003 000000000133ce94
0000c4002352a480
%l4-7: 0000000000000000 0000000000000002 0000000000000c0
0000000000000000
000002a1075c3970 genunix:ioctl+16c (3, 3103, fbf98, 3, 134d50, 0)
%l0-3: 0000c40040e0a50 000000000000c6d3 0000000000000003
0000030000002000
%l4-7: 0000000000000003 0000000000000004 0000000000000000
0000000000000000
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

### Oracle Solaris 10: Memory DR Operations Hang When Reducing Memory Below Four Gbytes

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

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

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

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

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

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

For example:

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

```
ldm rm-vcpu 760 ldg0
Request to remove cpu(s) sent, but no valid response received
VCPU(s) will remain allocated to the domain, but might
not be available to the guest OS
Resource removal failed
```

```
ldm set-vcpu 1 ldg0
Busy executing earlier command; please try again later.
Unable to remove the requested VCPUs from domain ldg0
Resource modification failed
```

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

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

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

Note that 760 exceeds the recommended maximum.

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

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

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

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

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

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

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

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

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

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

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

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

**Workaround:** None.

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

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

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

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

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

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

**Workaround:** Perform the following steps:

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

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

For information about rebuilding a domain configuration that uses virtual functions, see [“`ldm init-system` Command Might Not Correctly Restore a Domain Configuration on Which Physical I/O Changes Have Been Made” on page 52](#).

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

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

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

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

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

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

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

- A slot has been removed from a bus that is still owned by the primary domain.

- A virtual function has been created from a physical function that is owned by the primary domain.
- A virtual function has been assigned to the primary domain, to other guest domains, or to both.
- A root complex has been removed from the primary domain and assigned to a guest domain, and that root complex is used as the basis for further I/O virtualization operations.  
In other words, you created a non-primary root domain and performed any of the previous operations.

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

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

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

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

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

## Control Domain Requires the Lowest Core in the System

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

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

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

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

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

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

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

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

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

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

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

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

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

## Trying to Connect to Guest Domain Console While It Is Being Bound Might Cause Input to be Blocked

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**Workaround:** Perform the following steps:

1. Create an XML file for the primary domain.

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

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

```
ldm ls-constraints -x domain-name[,domain-name][,...] > 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. Apply the XML configuration to the guest domain or domains.

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

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

**Bug ID 15750727:** A system might panic when you reboot a primary domain that has a very large number of virtual functions assigned to it.

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

- Decrease the virtual function number to reduce the number of failed virtual functions. This change might keep the chip responsive.
- Create more Interrupt Resource Management (IRM) pools for the `ixgbe` virtual function because only one IRM pool is created by default for all the `ixgbe` virtual functions on the system.

## Partial Core primary Fails to Permit Whole-Core DR Transitions

**Bug ID 15748348:** When the primary domain shares the lowest physical core (usually 0) with another domain, attempts to set the whole-core constraint for the primary domain fail.

**Workaround:** Perform the following steps:

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

```
ldm list -o cpu
```

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

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

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

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

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

This problem has been reported only on a few systems.

**Workaround:** The status of the PCIe slots and the virtual functions is automatically updated after the information is received from the Logical Domains agent.

### **Removing a Large Number of CPUs From a Domain Might Fail**

The following bugs describe failures that might occur when removing a large number of CPUs from a domain.

- **Control domain.**

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

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

If you make a mistake while removing CPU resources, do *not* attempt another request to remove CPUs while the control domain is still in a delayed reconfiguration. If you do so, the commands will fail (see [“Only One CPU Configuration Operation Is Permitted to Be Performed During a Delayed Reconfiguration”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#)). Instead, undo the delayed reconfiguration operation by using the `ldm cancel-reconf` command, and start over.

3. Reboot the control domain.

- **Guest domain.**

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

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

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

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

**Bug ID 15721872:** You cannot use Oracle Solaris hot-plug operations to hot-remove a PCIe endpoint device after that device is removed from the primary domain by using the `ldm rm-io` command. For information about replacing or removing a PCIe endpoint device, see [“Making PCIe Hardware Changes”](#) in [“Oracle VM Server for SPARC 3.2 Administration Guide”](#).

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

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

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

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

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

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

**Bug ID 15702475:** A `No` response message might appear in the Oracle VM Server for SPARC log when a loaded domain's DRM policy expires after the CPU count has been substantially reduced. The `ldm list` output shows that more CPU resources are allocated to the domain than is shown in the `psrinfo` output.

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

## Virtual CPU Timeout Failures During DR

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

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

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

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

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

**Workaround:** Ignore this message.

## SPARC T3-1: Issue With Disks That Are Accessible Through Multiple Direct I/O Paths

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

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

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

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

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

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

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

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

## Guest Domains Running Oracle Solaris 10: Memory DR Removal Operations With Multiple Plumbed NIU `nxge` Instances Can Hang Indefinitely and Never Complete

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

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

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

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

```
ldm start-reconf primary
```

2. Assign the desired amount of memory to the domain.
3. Reboot the primary domain.

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

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

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

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

## An In-Use MAC Address Can Be Reassigned

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

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

### **ldmconfig Cannot Create a Domain Configuration on the SP**

**Bug ID 15654965:** The `ldmconfig` script cannot properly create a stored domain configuration on the service processor (SP).

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

1. Add the configuration to the SP.

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

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

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

3. Power cycle the system.

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

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

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

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

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

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

**Bug ID 15600969:** If all the hardware cryptographic units are dynamically removed from a running domain, the cryptographic framework fails to seamlessly switch to the software cryptographic providers, and kills all the `ssh` connections.

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

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

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

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

**Bug ID 15597025:** When you run the `ldm ls-io -l` command on a system that has a PCI Express Dual 10-Gigabit Ethernet Fiber card (X1027A-Z) installed, the output might show the following:

```
primary# ldm ls-io -l
...
pci@500/pci@0/pci@c PCIE5 OCC primary
network@0
network@0,1
ethernet
ethernet
```

The output shows four subdevices even though this Ethernet card has only two ports. This anomaly occurs because this card has four PCI functions. Two of these functions are disabled internally and appear as ethernet in the `ldm ls-io -l` output.

**Workaround:** You can ignore the ethernet entries in the `ldm ls-io -l` output.

### ldm Commands Are Slow to Respond When Several Domains Are Booting

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

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

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

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

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

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

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

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

For example:

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

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

```
ldm list -o network primary
NAME
primary

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

VSW
NAME MAC NET-DEV DEVICE DEFAULT-VLAN-ID PVID VID MTU MODE
```

```
primary-vsw0 00:14:4f:f9:86:f3 nxge0 switch@0 1 1 1500
```

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

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

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

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

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



---

**Caution** - Changes should not be made to `/etc/path_to_inst` without careful consideration.

---

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

**Bug ID 15555509:** When Logical Domains is configured on a system and you add another XAUI network card, the card is not visible after the machine has undergone a power cycle.

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

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

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

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

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

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

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

3. Perform a reconfiguration reboot of the control domain.

```
reboot -- -r
```

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

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

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

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

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

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

---

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

---

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

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

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

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

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

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

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

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

```
svcadm disable nis/client
```

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

**Bug ID 15511551:** Sometimes, executing the uadmin 1 0 command from the command line of a Logical Domains system does not leave the system at the ok prompt after the subsequent reset.

This incorrect behavior is seen only when the Logical Domains variable `auto-reboot?` is set to `true`. If `auto-reboot?` is set to `false`, the expected behavior occurs.

**Workaround:** Use this command instead:

```
uadmin 2 0
```

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

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

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

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

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

**Bug ID 15469227:** The `scadm` command on a control domain running at least the Oracle Solaris 10 5/08 OS can hang following an SC reset. The system is unable to properly re-establish a connection following an SC reset.

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

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

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

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

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

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

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

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

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

**Bug ID 15387338:** This issue is summarized in “[Logical Domains Variable Persistence](#)” in “[Oracle VM Server for SPARC 3.2 Administration Guide](#)” and affects only the control domain.

## Cannot Set Security Keys With Logical Domains Running

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

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

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

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

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

```
ldm stop-domain -f domain-name
```

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

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

## Resolved Issues

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.2 software release:

|          |                                                                                                            |
|----------|------------------------------------------------------------------------------------------------------------|
| 15486056 | Support migration using certificates                                                                       |
| 15573490 | Virtual switch should allow dynamic change of physical net-dev                                             |
| 15584929 | Add Logical Domains Manager support to config network bandwidth limit                                      |
| 15591769 | Virtual disk multipathing should provide options to manage paths                                           |
| 15594430 | Migration hvctl version warning message format error                                                       |
| 15637042 | Enhancements to the "ldm -V" output                                                                        |
| 15662016 | Remove UltraSPARC T1 support from ldmd                                                                     |
| 15677359 | Functional 64-bit Logical Domains Manager build needed                                                     |
| 15703117 | 'total_mem' and 'migrated_mem' fields in mig_cache_t are now unused                                        |
| 15716703 | Explicitly reject a live migration of a domain in OBP/KMDB                                                 |
| 15726854 | ldmd executes numerous misaligned loads during a migration                                                 |
| 15732015 | Upgrade compiler to Oracle Solaris Studio 12.3                                                             |
| 15732017 | Further optimize compilation of Logical Domains Manager                                                    |
| 15738764 | VM-API support for Oracle VM Server for SPARC                                                              |
| 15740276 | Make cpustat (perf counters) accessible from guest domains                                                 |
| 15748028 | ldmd should send an XML event when a OBP variable or VM-API pr                                             |
| 15770107 | Remove support for uncooperative migration                                                                 |
| 15776195 | Domain Dependencies Observability CLI and XML                                                              |
| 15781723 | Logical domain - virtual disk multipathing                                                                 |
| 15787684 | Logical Domains Manager should not roll back a memory DR operation while Solaris is processing the request |

---

|          |                                                                                                 |
|----------|-------------------------------------------------------------------------------------------------|
| 15798213 | Problem in MGMTLDMGR/LDOMMGR-MGMT                                                               |
| 15800189 | ldmp2v collect adds "s2" to end of slice information                                            |
| 15814176 | Need a way to create and destroy multiple virtual functions in one command                      |
| 15818139 | Core dump executing strcmp() in find_ldom_by_name()                                             |
| 15826354 | DR of large number of CPUs fails, but operation continues in guest                              |
| 15888594 | PAPSAT phase 1 tasks - ease of use, better graphs, and twiki cleanup                            |
| 15997828 | Modify LDoms-PM code to handle new compact PRI format                                           |
| 16164358 | Virtual function properties are missing from some ldm ls variants output.                       |
| 16368180 | Logical Domains Manager should provide an option to disable MAC addresses reallocation          |
| 16382628 | Logical Domains Manager vnet properties to be extended to cover important VNIC features         |
| 16404115 | Inconsistent state of guest set with max-ipc/whole-cores after doing DR(drd dis)                |
| 16494899 | Received invalid physical address during migration of domain                                    |
| 16618099 | XMLv3: Bindings: update physio section with new 'evacuated' property                            |
| 16908662 | /usr/sbin/ldmp2v[6]: syntax error at line 6 : `"' unmatched                                     |
| 16915664 | Assertion failed: ((chan-pending.msgs).count) == 0, file ds_comm.c, line 214                    |
| 16922112 | Logical Domains Manager core dump if disabled ldmd when there is a domain in a transition state |
| 16958880 | ldm stop-domain -t doesn't work as the man page explains                                        |
| 17012259 | Dynamic add/remove of nonexistent virtual functions fails with incorrect message.               |
| 17026219 | ldmp2v convert failure causes upgrade loop                                                      |
| 17043095 | null-pointer-deref in parse_mac_addr_list of iov_cmds.c:192                                     |

|          |                                                                                     |
|----------|-------------------------------------------------------------------------------------|
| 17043143 | use-after-free in client_loop of clients.c:735                                      |
| 17051532 | Removal of virtual function or PCIE device doesn't update autosave                  |
| 17069240 | Logical Domains Manager should support parallel MAC collision detection             |
| 17082986 | Saving spconfig via ldm add-spconfig -r returns out of memory error                 |
| 17179630 | LDoms support for Fibre Channel SR-IOV                                              |
| 17188920 | ldmpower's options suppress and timestamps are displaying nothing                   |
| 17205290 | Logical Domains Manager support to allocate memory for EQCB pool on M7/T7 platforms |
| 17215630 | Support for Fujitsu Board DR                                                        |
| 17235793 | PM Affinity CPU feature                                                             |
| 17285385 | Need to enable NPRD support for Fujitsu M10                                         |
| 17285811 | Migration bind invalid memory layout, panic on boot                                 |
| 17289361 | Allow -disable-migration-ssl in non-debug too                                       |
| 17296535 | Live migration needs configurable timeout value (and increased for debug builds)    |
| 17335156 | Integrate Fujitsu Board DR Support                                                  |
| 17343601 | Logical Domains Manager support to dynamically add/remove a PCIE bus to a domain    |
| 17344241 | Logical Domains Manager support for MR interrupt virtualization                     |
| 17344288 | Logical Domains Manager to restrict RKMVS root complexes to only primary domain     |
| 17355211 | Disable Direct I/O feature *ONLY* on M7/T7 platforms                                |
| 17362211 | Avoiding transfer of free list significantly shortens migration time (ldmd)         |
| 17362601 | Bring back "ldm ls -o" abbreviation support for "mem" and "net"                     |
| 17381876 | ldmp2v should use LC_ALL=C instead of LANG=C                                        |
| 17387444 | VM3.1 ldomVnetChange trap is not sent when alt-mac-addr is changed.                 |

---

|          |                                                                                                 |
|----------|-------------------------------------------------------------------------------------------------|
| 17391739 | SRIOV: destroy of virtual functions with INV stat shows: Attempt to online pf in primary failed |
| 17397557 | ldm migrate and beadm commands do not work together                                             |
| 17404856 | S10 does not support strnlen used in ds_vmapi                                                   |
| 17415778 | Coherency Link and shared policy reporting broken in PAPSAT                                     |
| 17418576 | ldmd does not return error exit code on ldm stop failure                                        |
| 17430911 | Oracle VM Server for SPARC doesn't honor Solaris policy on non-PAD S11 guests                   |
| 17431042 | Virtual LANs freed twice during vsw_add                                                         |
| 17432383 | Logical Domains Manager CLI to report all NICs and their properties from a given domain         |
| 17449926 | ldm add-domain -i ldom101.xml coredumps if file path is protected                               |
| 17463722 | New LDoms Domain Service to support net physdev listing and select linkprop mgmt                |
| 17464887 | Compression stream should be reused to avoid repeated memory allocation                         |
| 17468716 | Shared libraries should be built with position independent code (PIC)                           |
| 17502519 | ldm bind can't find a free VCC port for guest even if one is available                          |
| 17511365 | Promote Fujitsu Board DR development feature to product feature                                 |
| 17515863 | Add SPARC-M7 to 'generic' and 'migration-class1' cpu-arch lists                                 |
| 17556968 | Modify compact PRI parsing to handle SCC, Quadrant relationship for T7                          |
| 17561541 | ERROR: /pci@580: Invalid hypervisor argument(s). function: b4                                   |
| 17563510 | shutdown-group=0 for guest domain on Oracle systems                                             |
| 17596585 | ldmpower does not accommodate guest domain names longer than 13 characters                      |
| 17606070 | Poor affinity when set-mem done before set-core in delayed reconfig                             |

|          |                                                                                              |
|----------|----------------------------------------------------------------------------------------------|
| 17615719 | Logical Domains Manager should propagate all hv_mblock types                                 |
| 17621771 | OS panic after Board DR remaps degraded memory                                               |
| 17622266 | Need ability on debug Logical Domains Manager to specify memory RA addresses for DR commands |
| 17629136 | Two minor changes to LDoms parsing behavior                                                  |
| 17630764 | Enable dynamic bus feature based on existence of diov_rc_capability property                 |
| 17635306 | HV local memory not assigned if BB added without binding resources                           |
| 17636498 | On M7/T7, ldmpower calculation of normalized power should use SCC power info                 |
| 17663828 | Blacklist memory evacuation done on entire mblock instead of given memory range              |
| 17694771 | After migration VNET won't init: exceeded number of permitted handshake attempts             |
| 17707801 | Hypervisor needs to be aware of guest domain cpu-arch=sparc64-class1                         |
| 17709144 | Logical Domains Manager should support SP Domain Services over EoUSB                         |
| 17731533 | Memory allocations for Coprocessor in HVMD                                                   |
| 17742095 | sparc64-class1 migration failed since generic CPU version not set                            |
| 17759856 | Logical Domains Manager support to facilitate HV changes to track dirty pages in SPARC-M7    |
| 17761714 | ldm add-vcpu fails with Power Management Failure                                             |
| 17777004 | Remove degraded memory prior to Board DR/live migration                                      |
| 17777212 | Board DR core dump on attempt to remove unassigned memory region                             |
| 17785688 | Logical Domains Manager reports no ongoing DR when a memory DR op is stuck in a domain.      |
| 17796639 | ldmd core dump when attempting to bind domain in BINDING/ UNBINDING state                    |

---

|          |                                                                                      |
|----------|--------------------------------------------------------------------------------------|
| 17797079 | vnet_del_alt_mac_addrs() print a warning msg to SMF log for every failed match       |
| 17805392 | SPARC T5-1B S12 ldmd uses freed memory for ldm power handshake                       |
| 17818767 | ldmp2v Upgrade option not presented due to invalid /etc/vfstab                       |
| 17825714 | CPU utilization value too large for buffer length                                    |
| 17837437 | Oracle VM Server for SPARC core dump with ldm cancel-reconf                          |
| 17865561 | "ldm create-vf prop1=val1 prop2=val2 VF" only sets the first property (prop1)        |
| 17889357 | ldm set-vnet not updating vnet_node alt-mac-addr list for dup mac check              |
| 17898620 | [UPLSA] ldmpower(1M) with wrong options dumped core                                  |
| 17974700 | ldmd core dump on S10 during Board DR                                                |
| 17975946 | Physical function is being identified as null device when attempting to destroy VFs. |
| 18007132 | Alternative compression algorithms can reduce live migration time dramatically       |
| 18030537 | Support one system-wide LDC pool                                                     |
| 18045093 | Oracle VM Server for SPARC core dump on set-mem/rm-mem of S10 domain                 |
| 18047658 | Improve methods to diagnose ldmd fatal errors                                        |
| 18050132 | PRI update timeout during Board DR operation on M10                                  |
| 18066434 | Assertion failure on PRI update during Board DR                                      |
| 18106566 | ldmd should not allow pvid to be set on unsupported sxge hardware                    |
| 18112718 | Logical Domains Manager hits assert on cancel-reconf in pcie_unassign_all_devs       |
| 18116078 | Forward port 18115873 to 3.2 - migration of guest running pre s10u11 OS from 3.0     |
| 18123507 | Non related DR logs in ldmd for M5/M6                                                |

|          |                                                                                   |
|----------|-----------------------------------------------------------------------------------|
| 18169047 | Empty EMS slot shows UNK status following dynamic bus remove/add operations.      |
| 18169793 | ldomsmanager package should remove restart_fmri=svc:/system/rbac:default          |
| 18205211 | Miscellaneous hvctl cleanup, print response codes, static DTrace probe for errors |
| 18225554 | Logical Domains Manager package descriptions need some updating                   |
| 18236261 | Disabled policy doesn't force cpus to highest power state                         |
| 18269210 | After delay-reconf, reboot NPRD caused ldmd core dumped                           |
| 18282403 | Policy int should be queried only if the policy budget str can't be obtained      |
| 18298193 | Live migration should read memory pages out of HV in parallel (ldmd)              |
| 18308270 | cppcheck reports a number of issues in the LDOMs codebase.                        |
| 18320741 | migrate_init() should use strlen on src_hostname rather than unsafe strlen        |
| 18328493 | Fix handling of \$ROOT and remove dependency on ws(1)                             |
| 18338247 | Power management bug fixes and enhancements found during PM Quadrant development  |
| 18375366 | Fix some Uninitialised Memory Errors reported by Parfait 1.4 RR                   |
| 18399030 | New ds_netsvc domain service and vnet linkprop support needed                     |
| 18402614 | Support Tahoe static coprocessor frags                                            |
| 18477335 | Hypervisor Abort occurred on deleteboard unbind=resource                          |
| 18479243 | ldmp2v_prepare failing on large partition                                         |
| 18488324 | Improve performance of pri_reinit_cpus()                                          |
| 18511736 | Memory leak in p2v/disk.c                                                         |
| 18511763 | Memory leak LDOMS 3.2. File iov_cmds.c                                            |
| 18531327 | Fujitsu CMI OVM support                                                           |

---

|          |                                                                                           |
|----------|-------------------------------------------------------------------------------------------|
| 18538035 | ldm destroy triggers assertion failure: !vfs-bound, file cons_eng.c, line 721,            |
| 18558868 | Tidy mig_common.h and move definitions to mig.h or new mig_debug.h where applic.          |
| 18594819 | Enable dynamic PCIe slots on Fujitsu M10                                                  |
| 18595023 | SR-IOV cannot be used on PCI-BOX SLOT#4 or higher                                         |
| 18595954 | ldmd core dumps on add-io operations.                                                     |
| 18600414 | ldmd impacting removal of cpus on m632                                                    |
| 18603309 | NAME TYPE are conjuncted in "ldm ls-io" output when IO NAME is too long                   |
| 18613122 | ds_service_routine() and ds_channel_service_routine() are equivalent                      |
| 18639528 | Isolating a domain prior to suspension causes problems for KZ suspend/migration           |
| 18664262 | cpu_del() needs comments                                                                  |
| 18664570 | ldmd crashes with a SEGV on attempting VF recovery.                                       |
| 18665751 | Oracle VM Server for SPARC workaround to fix DRM on Fujitsu M10 running XCP2210           |
| 18672454 | Logical Domains Manager fails to clean up data structures if LDoms agent returns an error |
| 18673124 | Live Repair Resource Group Observability CLI and XML                                      |
| 18746688 | Recreating autosave guest MDs significantly degrades CLI responsiveness                   |
| 18751174 | Misc minor fixes to migration code                                                        |
| 18770805 | Recovery mode fails with ldmd in maintenance mode when VSW net-dev is missing.            |
| 18778422 | prtconf delays cause ldmd to drop to maintenance                                          |
| 18795232 | ldmd crashes on live migration during Board DR                                            |
| 18795371 | Could not obtain PRI on large configuration Fujitsu M10                                   |

|          |                                                                                 |
|----------|---------------------------------------------------------------------------------|
| 18803774 | CMU DR failed: WARNING: Unexpected status: DR_FAILURE by GM                     |
| 18895023 | problem in MGMTLDMGR/LDOMMGR-MGMT                                               |
| 18915166 | ldom duplicate mac address detection doesn't check alt-mac-addr on domain start |
| 18951249 | ldmd_start's create_key() needs to be more robust                               |
| 18952009 | Port Bug 18726175 - ldmd aborts when adding memory to guest domain to 3.2       |
| 18997349 | Resource group 'core-remove' and 'core-move' commands                           |
| 18997448 | Resource group 'memory-remove' and 'memory-move' commands                       |
| 19015999 | Assert failure in suspended reconfig: 'hv_susrec_reply == 0'                    |
| 19027934 | add-domain and set-domain usage info has excess information and missing braces  |
| 19051705 | Logical Domains Manager support for PVLAN                                       |
| 19072933 | Use RSA-2048 and SHA-256 for self-signed certificates                           |
| 19074284 | Use /dev/random not urandom when generating keys/certs                          |
| 19079057 | init-system dumps core when the input XML doesn't have any services config      |
| 19079764 | ldmd core dump at cpu_del+0xce4                                                 |
| 19172925 | Provide debug command to help stress test solaris freelist support              |
| 19172946 | Support freelist responses with non-zero RA offsets                             |
| 19173667 | ldm add-spconfig should also update autosave                                    |
| 19263879 | Enable 'dr-pdom' domain service on Oracle platforms                             |
| 19263922 | Add support for hold/release of a strand to 'dr-pdom' DS                        |
| 19310514 | Added resources not assigned when restoring degraded primary failed             |
| 19310540 | addboard to factory-default does not correctly assign CPU resources             |
| 19310550 | ldmd service in maintenance state if large hvdump data                          |

---

|          |                                                                                 |
|----------|---------------------------------------------------------------------------------|
| 19310579 | Faulty CPU core not autoreplaced after PPAR reset                               |
| 19310587 | Guest domain state left "suspending" on domain suspend failure                  |
| 19322662 | Problem in MGMTDS/DS-MGMT                                                       |
| 19324709 | Virtual function info missing in ls-bindings and ls -o physio                   |
| 19334816 | Fix incorrect format strings hidden behind gettext()                            |
| 19345470 | Fatal error: The PRI does not contain a singleton cod-license node              |
| 19356766 | set-vcpu error message when running with dynamic pools needs to be cleaned up   |
| 19368156 | Clean up dr_mem_unconfigure_range() for Live Repair                             |
| 19372065 | Adding a vnet with an invalid vsw name fails silently                           |
| 19379138 | In multi-master dependency cfg, slaves do not take policy of 1st master to fail |
| 19379365 | 256M alignment needed for M10-4S PPAR DR                                        |
| 19422259 | Frag remap after live migration may crash ldmd                                  |
| 19424242 | ldmd crashes repeatedly on M10 if I/O domain degraded                           |
| 19424359 | hvdump properties lost during recovery mode                                     |
| 19429567 | ldmd should check for connections periodically, case failed initial connection  |
| 19430009 | Logical Domains Manager crashes when processing a binary MD containing errors   |
| 19450132 | Physical memory constraints lost if all memory degraded                         |
| 19456310 | Hypervisor abort on set-core of control domain                                  |
| 19484790 | ldm add-vswitch error message could be more descriptive – invalid device        |
| 19503255 | Removing vnets or disks on guest causes ldmd crash and restart.                 |
| 19511743 | Core API group at version 1.3 only available on newer firmware (9.2.X)          |

|          |                                                                                             |
|----------|---------------------------------------------------------------------------------------------|
| 19513561 | Live migration failure on M10 crashed ldmd                                                  |
| 19519214 | Domain migration should support FIPS 140-2 mode                                             |
| 19521069 | Forwardport 19480835 Very large LDC counts can result in Solaris issues in guest            |
| 19614734 | "ls-netdev" cmd should have an option to list the possible backend devs for vsw             |
| 19674029 | A move-mem _sys_ should fail if board has active IOS segments                               |
| 19676615 | vds_find_volume_by_name() needs to be more careful                                          |
| 19781715 | deleteboard unbind=resource fails repeatedly fails                                          |
| 19793172 | ldmd au_write negative audit descriptor: Bad file number                                    |
| 19794564 | Logical Domains Manager requires write authorization to register for XMPP config events     |
| 19809799 | Secondary VCC configurations not properly restored with an init-system restore              |
| 19826474 | Oracle VM Server for SPARC 3.2 B34 deleteboard fails if all board memory resources are free |
| 19864344 | Explicitly disable SSLv2 and SSLv3 to avoid POODLE (CVE-2014-3566)                          |
| 20064006 | FC-SRIOV: stale data shown from "ldm ls-io" after having done a destroy-vf op               |
| 20157326 | fma-io-domain-service returns an unassigned status                                          |
| 20177235 | Logical Domains Manager assumes the 'factory-default' SP config timestamp is zero           |
| 20319742 | Update Logical Domains Manager package copyright to current year                            |