

# Oracle<sup>®</sup> VM Server for SPARC 3.4 Release Notes

ORACLE<sup>®</sup>

Part No: E69555  
June 2016



**Part No: E69555**

Copyright © 2007, 2016, 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.

**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.

**Référence: E69555**

Copyright © 2007, 2016, 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, accorder de licence, 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 livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel 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 un risque de dommages corporels. Si vous utilisez ce logiciel ou ce 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 des applications dangereuses.

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 de 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.

**Accès aux services de support Oracle**

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 3.4 Release Notes** ..... 9
  - What's New in This Release ..... 10
  - Features That Depend on System Firmware, the Oracle Solaris OS, or Both ..... 12
  - Oracle VM Server for SPARC 3.4 System Requirements ..... 12
  - Deprecated and Removed Oracle VM Server for SPARC Features ..... 13
  - Known Issues ..... 13
    - Migration Issues ..... 13
    - Bugs Affecting the Oracle VM Server for SPARC Software ..... 22
    - Documentation Issues ..... 40
  - Resolved Issues ..... 40



## Using This Documentation

---

- **Overview** – Provides late-breaking information about the Oracle VM Server for SPARC 3.4 software, such as changes for this release and known bugs that affect the software.
- **Audience** – System administrators who manage virtualization on SPARC servers.
- **Required knowledge** – System administrators with working knowledge of UNIX systems and the Oracle Solaris operating system (Oracle Solaris OS).

## Product Documentation Library

Documentation and resources for this product and related products are available at <http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html>.

## Feedback

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



## Oracle VM Server for SPARC 3.4 Release Notes

---

These release notes include information about issues that you might encounter if you are running the fully qualified versions of the Oracle Solaris OS, system firmware for a supported platform, and the Oracle VM Server for SPARC 3.4 software. If you are not running these qualified versions, you might encounter a larger set of issues.

---

**Note** - Ensure that you install and run the Oracle VM Server for SPARC 3.4 software with the fully qualified system firmware versions on the supported hardware platforms. All domains on the system must run the latest Support Repository Update (SRU) of an Oracle Solaris 11 OS or the latest patch for the Oracle Solaris 10 1/13 OS.

These release notes might include some known issues that exist with older versions of the software.

---

For information about the supported hardware and fully qualified software and system firmware, see [Chapter 1, “System Requirements” in \*Oracle VM Server for SPARC 3.4 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 3.4 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 SPARC T-series servers and SPARC 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 servers 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 3.4 Installation Guide*. 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” in *Oracle VM Server for SPARC 3.4 Release Notes* 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 \(http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html\)](http://www.oracle.com/technetwork/server-storage/vm/documentation/sparc-whatsnew-330281.html).

The major changes for the Oracle VM Server for SPARC 3.4 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 3.4 Installation Guide*.

- Added support for the SPARC S7 series servers. See “Supported Platforms” in *Oracle VM Server for SPARC 3.4 Installation Guide*.
- Enhanced the virtual SCSI HBA feature to enable I/O multipathing in the service domain. See Chapter 12, “Using Virtual SCSI Host Bus Adapters” in *Oracle VM Server for SPARC 3.4 Administration Guide*.  
  
This feature is also available on UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers. However, these older servers do not support booting from a virtual SCSI HBA.
- Added MIB support for the virtual SCSI HBA-related `ldm` list commands. See *Oracle VM Server for SPARC 3.4 Management Information Base User's Guide*.
- Enabled a trusted host to create virtual devices such as VLANs and VNICs by permitting guest-specified MAC addresses and VLAN IDs. See “Using Trusted Virtual Networks” in *Oracle VM Server for SPARC 3.4 Administration Guide*.
- Extended the verified boot feature to include guest domains. This feature enables you to validate the signatures of SPARC boot blocks and Oracle Solaris kernel modules. See “Using Verified Boot” in *Oracle VM Server for SPARC 3.4 Administration Guide*.
- Changed the default `inter-vnet-link` property value to `auto`. The default behavior ensures that inter-vnet LDC channels are enabled unless more than eight virtual networks are associated with a single virtual switch. This property enforces the switching off of inter-vnet-links if the number exceeds the configured value, which is eight by default. See “Inter-

[Vnet LDC Channels](#)” in *Oracle VM Server for SPARC 3.4 Administration Guide* and the [ldm\(1M\)](#) man page.

This feature is also available on UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers.

- Enhanced property modification methods for the `ldm add-vnet`, `ldm set-vnet`, `ldm add-vsw`, and `set-vsw` commands. See the [ldm\(1M\)](#) man page.

This feature is also available on UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers.

- Expanded the number of LDC channels to 4096 on several platforms.
- Added support for the dynamic modification of device-specific properties on InfiniBand SR-IOV devices. Use the `ldm set-io property=value InfiniBand-VF-or-PF` command.
- Changed the default `linkprop` property value to `phys-state`. See the [ldm\(1M\)](#) man page.

This feature is also available on UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers.

- Enhanced the `ldm list-bindings` command to show details about virtual network and virtual switch devices. See the [ldm\(1M\)](#) man page.

This feature is also available on UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers.

- Enhanced the Oracle VM Server for SPARC template utilities and add the `ovmtadm` command. See the [ovmtadm\(1M\)](#) man page.

This feature is also available on UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers.

- Restricted the XML interface to only support at least version 1.2 of the Transport Layer Security (TLS) protocol. Support for deprecated protocols such as SSLv3, TLSv1.0, and TLSv1.1 has ended.

This feature is also available on UltraSPARC T2, UltraSPARC T2 Plus and SPARC T3 servers.

- Added support for physical partition dynamic reconfiguration policy options for Fujitsu M10 servers. See *Fujitsu M10/SPARC M10 Systems System Operation and Administration Guide*.

This feature is supported only on Fujitsu M10 servers.

- Bug fixes.

## Features That Depend on System Firmware, the Oracle Solaris OS, or Both

Some of the Oracle VM Server for SPARC 3.4 features are available only when the fully qualified system firmware and Oracle Solaris OS are installed. Note that the control domain must already be running the fully qualified Oracle Solaris OS.

For information about the supported hardware, system firmware, and Oracle Solaris OS, see [Chapter 1, “System Requirements” in \*Oracle VM Server for SPARC 3.4 Installation Guide\*](#). For Fujitsu M10 servers, see the latest *Fujitsu M10/SPARC M10 Systems Product Notes*.

For Oracle VM Server for SPARC 3.4, all new features are available even if the system does not run the fully qualified system firmware.

If a guest domain, an I/O domain, or a root domain does not run the fully qualified Oracle Solaris OS version, the following Oracle VM Server for SPARC 3.4 features are not available:

- Verified boot
- Trusted virtual networks
- Increased number of LDC channels
- Dynamic modification of device-specific properties on InfiniBand SR-IOV devices
- `linkprop=phys` - state by default
- Physical partition dynamic reconfiguration policy options for Fujitsu M10 servers requires at least the Oracle Solaris 11.3 OS

## Oracle VM Server for SPARC 3.4 System Requirements

You can find information about the recommended and minimum software component versions to use with the Oracle VM Server for SPARC 3.4 release in [Chapter 1, “System Requirements” in \*Oracle VM Server for SPARC 3.4 Installation Guide\*](#).

## Deprecated and Removed Oracle VM Server for SPARC Features

The following previously deprecated Oracle VM Server for SPARC features have been removed from this release of the software:

- Setting the `threading` property to manage single-thread workloads on SPARC T4 servers. 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. This feature has been removed from the Oracle VM Server for SPARC 3.3 release.

Before you install or upgrade to the Oracle VM Server for SPARC 3.4 software, ensure that all domains have the `threading` property set to `max-throughput`. See [“How to Remove the threading Property”](#) in *Oracle VM Server for SPARC 3.4 Installation Guide* and [“Upgrading the Oracle Solaris OS”](#) in *Oracle VM Server for SPARC 3.4 Installation Guide*.

- The Netra Data Plane software suite, which includes the `vdpc` and `vdpcs` virtual devices, is no longer supported by Oracle VM Server for SPARC.
- Using the Hybrid I/O feature is deprecated in favor of using the single-root I/O virtualization (SR-IOV) feature.
- The Logical Domains Manager auditing implementation has been removed pending replacement.

## Known Issues

This section contains general issues and specific bugs concerning the Oracle VM Server for SPARC 3.4 software.

### Migration Issues

#### **Inaccurate Unable to Send Suspend Request Error Reported During a Successful Domain Migration**

**Bug ID 23206413:** In rare circumstances, a successful domain migration reports the following error:

Unable to send suspend request to domain *domain-name*

This issue occurs when the Logical Domains Manager detects an error while suspending the domain, and the Logical Domains Manager is able to recover and completes the migration. The exit status of the command is 0, reflecting the successful migration.

**Workaround:** Because the migration completes successfully, you can ignore the error message.

## Migrating a Bound Domain With Many Virtual Devices Might Fail and Leave Two Bound Copies of the Domain

**Bug ID 23180427:** When migrating a bound domain that has a large number of virtual devices, the operation might fail with the following message in the SMF log:

```
warning: Timer expired: Failed to read feasibility response type (9) from target LDoms Manager
```

This failure indicates that the Logical Domains Manager running on the source machine timed out while waiting for the domain to be bound on the target machine. The chances of encountering this problem increases as the number of virtual devices in the migrating domain increases.

The timing of this failure results in a bound copy of the domain on both the source machine and the target machine. Do not start both copies of this domain. This action can cause data corruption because both domains reference the same virtual disk backends.

**Recovery:** After verifying that the copy of the migrated domain is correct on the target machine, manually unbind the copy of the domain on the source machine and destroy it.

## Migration Fails When the Target Machine Has Insufficient Free LDCs

**Bug ID 23031413:** When the target machine's control domain runs out of LDCs during a domain migration, the migration fails and the following message is written to the SMF log:

```
warning: Failed to read feasibility response type (5) from target LDoms Manager
```

This error is issued when the domain being migrated fails to bind on the target machine. Note that the bind operation might fail for other reasons on the target machine, as well.

**Workaround:** For the migration to succeed, the number of LDCs must be reduced either in the domain being migrated or in the control domain of the target machine. You can reduce the

---

number of LDCs by reducing the number of virtual devices being used by or being serviced by a domain. For more information about managing LDCs, see [“Using Logical Domain Channels” in Oracle VM Server for SPARC 3.4 Administration Guide](#).

## Domain Migration Is Supported Only With at Least TLS v1.2

**Bug ID 23026264:** Starting with Oracle VM Server for SPARC 3.4, the Logical Domains Manager only supports at least TLS v1.2 for secure domain migrations. If the peer of the migration is incapable of using TLS v1.2, the migration fails with the following error message:

```
Failed to establish connection with ldmd(1m) on target: target
Check that the 'ldmd' service is enabled on the target machine and
that the version supports Domain Migration. Check that the 'xmpc_enabled'
and 'incoming_migration_enabled' properties of the 'ldmd' service on
the target machine are set to 'true' using svccfg(1M).
```

Domain migration is supported only between two consecutive minor versions of the Oracle VM Server for SPARC software. This problem does not affect any of the supported combinations. However, Oracle VM Server for SPARC software running on the Oracle Solaris 10 OS is unable to use TLS v1.2 by default and is incompatible for domain migration with Oracle VM Server for SPARC 3.4.

---

**Note** - This is a generic error message that you might encounter in other circumstances including providing an incorrect password.

---

## boot-policy Property Value Is Not Preserved When a Guest Domain Is Migrated to an Older Oracle VM Server for SPARC Version and Is Later Migrated to Oracle VM Server for SPARC 3.4

**Bug ID 23025921:** The boot-policy property of a guest domain is not preserved when the guest domain is migrated to a system that runs an older version of the Logical Domains Manager and is later migrated to a system that runs Oracle VM Server for SPARC 3.4.

The Oracle VM Server for SPARC 3.4 software introduced the boot-policy property to support the verified boot feature. Older versions of the Oracle VM Server for SPARC software do not support this property, so the boot-policy property is dropped when a guest domain is migrated from a system that runs Oracle VM Server for SPARC 3.4 to a system that runs a version of Oracle VM Server for SPARC that is older than 3.4.

When the guest domain is later migrated to a system that runs Oracle VM Server for SPARC 3.4, the default `boot-policy` value of `warning` is applied to the migrated guest domain.

**Recovery:** After migrating the guest domain to a target system that runs Oracle VM Server for SPARC 3.4, manually set the `boot-policy` property to the desired value. Perform this step if the default value of `warning` is not appropriate.

1. Set the `boot-policy=none`.

```
primary# ldm set-domain boot-policy=none ldg1
```

2. Reboot the guest to make the new boot policy take effect.

## Kernel Zones Block Live Migration of Guest Domains

**Bug ID 21289174:** On a SPARC server, a running kernel zone within an Oracle VM Server for SPARC domain will block live migration of the guest domain. The following error message is shown:

```
Guest suspension 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
```

- Perform a live migration of the kernel zone to another system before migrating the guest domain.

See [Chapter 3, “Migrating an Oracle Solaris Kernel Zone”](#) in *Creating and Using Oracle Solaris Kernel Zones*.

## 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.

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

**Bug ID 17285751:** Migrating an Oracle Solaris 10 guest 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()`.

Note that this problem has been fixed in the Oracle Solaris 11.1 OS.

**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.

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

**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 server and an UltraSPARC T2 or SPARC T3 server.

**Workaround:** None.

## `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 an incorrect 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.

## **Guest Domain Panics When Running the `cputrack` Command During a Migration to a SPARC T4 Server**

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

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

## **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 servers.

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.

## **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 the migration succeeds, the unplanned restart of the migrated domain on the target machine means that its not 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.

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

**Bug ID 15696986:** If two `ldm migrate` commands are issued between the same two systems 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.

## 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 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 3.4 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.

Note that this problem occurs when the migrated domain is running an OS version older than Oracle Solaris 11.3.

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

```
# svcadm restart vntsd
```

---

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

---

## Cannot Migrate a Domain Between a System That Has EFI GPT Disk Labels and a System That Does Not Have EFI GPT Disk Labels

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

System firmware versions 8.4, 9.1, and XCP2230 introduced support for EFI GPT disk labels. By default, virtual disks that are installed when running at least the Oracle Solaris 11.1 OS on those systems have an EFI GPT disk label. You cannot read this disk label on older versions of firmware (such as 9.0.x, 8.3, 7.x, or XCP2221). This situation precludes you from performing a live migration or a cold migration to a system that runs a system firmware version without EFI GPT support. Note that a cold migration also fails in this situation, which is different than the previous limitations.

To determine whether your virtual disk has an EFI GPT disk label, run the `devinfo -i` command on the raw device. The following examples show whether the virtual disk has an SMI VTOC or an EFI GPT disk label.

- **SMI VTOC disk label.** When your virtual disk has an SMI VTOC, you can perform a migration to firmware regardless of whether it supports EFI.

This example indicates that the device has a VTOC label because the `devinfo -i` command reports device-specific information.

```
# devinfo -i /dev/rdisk/c2d0s2
/dev/rdisk/c2d0s2      0      0      73728  512    2
```

- **EFI GPT disk label.** When your virtual disk has an EFI GPT disk label, you can perform a migration only to firmware that has EFI support.

This example indicates that the device has an EFI GPT disk label because the `devinfo -i` command reports an error.

```
# devinfo -i /dev/rdisk/c1d0s0
devinfo: /dev/rdisk/c1d0s0: This operation is not supported on EFI
labeled devices
```

## 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.

### Bugs Affecting the Oracle VM Server for SPARC 3.4 Software

#### Support for Static Virtual Function Creation During Recovery Mode

**Bug ID 23205662:** Due to a limitation in the PSIF driver used by certain InfiniBand cards, the driver does not support dynamic IOV operations such as virtual function creation. This limitation results in recovery mode failing to recover non-primary root domains that have physical functions that use the PSIF driver. The physical functions never become ready to create virtual functions because of the lack of support for dynamic IOV operations.

**Workaround:** Do not create virtual functions on InfiniBand physical functions that use the PSIF driver in non-primary root domains.

#### I/O Domain Recovery Fails With Virtual Functions in an Invalid State

**Bug ID 23170671:** Sometimes virtual functions and physical functions remain in an invalid state after creating virtual functions. A domain that has such a virtual function assigned to it cannot be bound. If this issue occurs during recovery mode, any I/O domains that have virtual functions in the INV state are not recovered.

The `ldmd` log shows messages similar to the following for the `IOVFC.PF1` physical function:

```
Recreating VFs for PF /SYS/MB/PCIE2/IOVFC.PF0 in domain root_2
Recreating VFs for PF /SYS/MB/PCIE2/IOVFC.PF1 in domain root_2
Recreating VFs for PF /SYS/MB/NET2/IOVNET.PF0 in domain root_3
PF /SYS/MB/PCIE2/IOVFC.PF1 not ready (3)
PF /SYS/MB/PCIE2/IOVFC.PF1 not ready (3)
PF /SYS/MB/PCIE2/IOVFC.PF1 not ready (3)
PF /SYS/MB/PCIE2/IOVFC.PF1 not ready (3)
```

**Recovery:** If you notice this problem in time, you can restart the `ldmd` agent in the `root_2` domain to resolve this issue while recovery mode continues to retry the physical function. Restarting the agent enables the recovery of the I/O domains that use virtual functions of the physical function. If you do not notice this problem in time, the recovery operation continues but will not be able to recover the I/O domains that use those virtual functions.

## Oracle VM Server for SPARC MIB `ldomSPConfigTable` Does Not Show All SP Configurations

**Bug ID 23144895:** Oracle VM Server for SPARC MIB shows only the factory-default configuration for the service processor (SP) configuration table (`ldomSPConfigTable`).

**Workaround:** To show the complete list of the SP configurations on the system, use the `ldm list-spconfig` or the `list-spconfig XML` interface.

For example:

```
primary# ldm list-spconfig
factory-default [next poweron]
test_config
```

The XML `list-spconfig` responds as follows:

```
<cmd>
  <action>list-spconfig</action>
  <data version="3.0">
    <Envelope>
      <References/>
      <Section>
        <Item>
          <rasd:OtherResourceType>spconfig</rasd:OtherResourceType>
          <gprop:GenericProperty key="spconfig_name">factory-default</gprop:
GenericProperty>
          <gprop:GenericProperty key="spconfig_status">next</gprop:GenericProperty>
        </Item>
      </Section>
      <References/>
      <Section>
        <Item>
          <rasd:OtherResourceType>spconfig</rasd:OtherResourceType>
          <gprop:GenericProperty key="spconfig_name">test_config</gprop:GenericProperty>
        </Item>
      </Section>
    </Envelope>
  </data>
  ...
```

## ovmtlibrary Limits Disk Image File Name to 50 Characters

**Bug ID 23024583:** The `ovmtlibrary` command limits the disk image file name to 50 characters. The `ovmtlibrary` checks the `.ovf` file and compares the information in the `<ovf:References>` section with the actual file names of the decompressed disks.

An error is issued if the files are different or if the disk image file name is longer than 50 characters. For example:

```
# ovmtlibrary -c store -d "example" -q -o file:/template.ova -l /export/user1/
ovmtlibrary_example
event id is 3
ERROR: The actual disk image file name(s) or the actual number of disk
image(s) is different from OVF file: template.ovf
exit code: 1
```

The following example XML shows a disk image file name that is greater than 50 characters:

```
<ovf:References>
<ovf:File ovf:compression="gzip"
ovf:href="disk_image.ldoms3.4_build_s11_u3_sru06_rti_02_kz_40G.img.gz"
ovf:id="ldoms3" ovf:size="6687633773"/>
</ovf:References>
```

**Workaround:** Limit the length of disk image file names to fewer than 50 characters.

## ovmtcreate Creates a Corrupt Template if the Same vdsdev Backend File Name Is Found

**Bug ID 22919488:** The `ovmtcreate` command does not support the creation of templates from source domains where the `vdsdev` has the same name for more than one virtual disk in the same domain.

This problem is unlikely to occur because source domains that have multiple virtual disks typically have different backend devices and therefore different file names. However, if `ovmtdeploy` is used with a template that has been created from a source domain where the `vdsdev` has the same name for more than one virtual disk, `ovmtdeploy` fails with an error message. For example:

```
# ovmtdeploy -d ldg1 template.ova
ERROR: pigz:
//ldg1/resources/disk_image.ldoms3.4_build_s11_u3_sru05_rti_01_kz_36G.img.gz
does not exist -- skipping
FATAL: Failed to decompress disk image
```

**Workaround:** Specify different vdsdev backend file names for virtual disks that are contained in the same domain.

### **Virtual Network Devices Added to an Inactive Guest Domain Never Gets the Default linkprop Value**

**Bug ID 22842188:** For `linkprop=phys-state` to be supported on a virtual network device, the Logical Domains Manager must be able to validate that the virtual switch to which the virtual network device is attached has a physical NIC that backs the virtual switch.

The Oracle VM Server for SPARC `netsvc` agent must be running on the guest domain so that the virtual switch can be queried.

If the guest domain is not active and cannot communicate with the agent in the domain that has the virtual network device's virtual switch, the virtual network device does not have `linkprop=phys-state` set.

**Workaround:** Only set `linkprop=phys-state` when the domain is active.

### **`ldm set-vsw net-dev=` Fails When `linkprop=phys-state`**

**Bug ID 22828100:** If a virtual switch has attached virtual network devices that have `linkprop=phys-state`, the virtual switch to which they are attached must have a valid backing NIC device specified by the `net-dev` property. The `net-dev` property value must be the name of a valid network device.

If this action is performed using `net-dev=`, the virtual switch still shows `linkprop=phys-state` even though the `net-dev` property value is not a valid NIC device.

**Workaround:** First, remove all the virtual network devices that are attached to the virtual switch, and then remove the virtual switch. Then, re-create the virtual switch with a valid `net-dev` backing device, and then re-create all the virtual network devices.

### **A Domain That Has Socket Constraints Cannot Be Re-Created From an XML File**

**Bug ID 21616429:** The Oracle VM Server for SPARC 3.3 software introduced socket support for Fujitsu M10 servers only.

Software running on Oracle SPARC servers and Oracle VM Server for SPARC versions older than 3.3 cannot re-create a domain with socket constraints from an XML file.

Attempting to re-create a domain with socket constraints from an XML file with an older version of the Oracle VM Server for SPARC software or on an Oracle SPARC server fails with the following message:

```
primary# ldm add-domain -i ovm3.3_socket_ovm11.xml
socket not a known resource
```

If Oracle VM Server for SPARC 3.2 is running on a Fujitsu M10 server and you attempt to re-create a domain with socket constraints from an XML file, the command fails with various error messages, such as the following:

```
primary# ldm add-domain -i ovm3.3_socket_ovm11.xml
Unknown property: vcpus
```

```
primary# ldm add-domain -i ovm3.3_socket_ovm11.xml
perf-counters property not supported, platform does not have
performance register access capability, ignoring constraint setting.
```

**Workaround:** Edit the XML file to remove any sections that reference the socket resource type.

## Slow I/O on Virtual SCSI HBA Guest Domain When One of the Service Domains Is Down With a Virtual SCSI HBA Timeout Set

**Bug ID 21321166:** I/O throughput is sometimes slower when using a virtual SCSI HBA MPxIO path to an offline service domain.

**Workaround:** Disable the path to the offline service domain by using the `mpathadm disable path` command until the service domain is returned to service.

## Virtual SCSI HBA Does Not See Dynamic LUN Changes Without a Reboot

**Bug ID 21188211:** If LUNs are added to or removed from a virtual SAN after a virtual SCSI HBA is configured, the `ldm rescan-vhba` command sometimes does not show the new LUN view.

**Workaround:** Remove the virtual SCSI HBA and then re-add it. Check to see whether the LUNs are seen. If the removal and re-add operations are unsuccessful, you must reboot the guest domain.

### **mpathadm Shows Incorrect Path State Output for a Virtual SCSI HBA When a Fibre Channel Cable Is Pulled**

**Bug ID 20876502:** Pulling the SAN cable from a service domain that is part of a virtual SCSI HBA MPxIO guest domain configuration causes the Path State column of the `mpathadm` output to show incorrect values.

**Workaround:** Plug in the SAN cable and run the `ldm rescan-vhba` command for all the virtual SCSI HBAs to the service domain that has the cable attached. After performing this workaround, the guest domain should resume performing I/O operations.

### **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.

---

**Note** - This problem applies to UltraSPARC T2, UltraSPARC T2 Plus, SPARC T3, and SPARC T4 series servers. This problem also applies to SPARC T5, SPARC M5, and SPARC M6 series servers that run a system firmware version prior to 9.5.3.

---

**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.

### **Panic When Using the `ldm rm-io virtual-function` Command to MPxIO That Contains a Virtual SCSI HBA**

**Bug ID 20046234:** When a virtual SCSI HBA and a Fibre Channel SR-IOV device can view the same LUNs in a guest domain when MPxIO is enabled, a panic might occur. The panic occurs if the Fibre Channel SR-IOV card is removed from the guest domain and then re-added.

**Workaround:** Do not configure a guest domain with Fibre Channel SR-IOV and a virtual SCSI HBA when both have MPxIO enabled.

### **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.

## 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.

## Misleading Messages Shown for InfiniBand SR-IOV Remove Operations

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

Dynamic SR-IOV remove operations are not supported for InfiniBand devices.

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

- [“How to Remove an InfiniBand Virtual Function From an I/O Domain”](#) in *Oracle VM Server for SPARC 3.4 Administration Guide*
- [“How to Remove an InfiniBand Virtual Function From a Root Domain”](#) in *Oracle VM Server for SPARC 3.4 Administration Guide*

## Resilient I/O Domain Should Support PCI Device Configuration Changes After the Root Domain Is Rebooted

**Bug ID 16691046:** If virtual functions are assigned from the root domain, an I/O domain might fail to provide resiliency in the following hotplug situations:

- You add a root complex (PCIe bus) dynamically to the root domain, and then you create the virtual functions and assign them to the I/O domain.
- You hot-add an SR-IOV card to the root domain that owns the root complex, and then you create the virtual functions and assign them to the I/O domain.
- You replace or add any PCIe card to an empty slot (either through hotplug or when the root domain is down) on the root complex that is owned by the root domain. This root domain provides virtual functions from the root complex to the I/O domain.

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

- If the root complex already provides virtual functions to the I/O domain and you add, remove, or replace any PCIe card on that root complex (through hotplug or when the root domain is down), you must reboot both the root domain and the I/O domain.
- If the root complex does not have virtual functions currently assigned to the I/O domain and you add an SR-IOV card or any other PCIe card to the root complex, you must stop the root domain to add the PCIe card. After the root domain reboots, you can assign virtual functions from that root complex to the I/O domain.
- If you want to add a new PCIe bus to the root domain and then create and assign virtual functions from that bus to the I/O domain, perform one of the following steps and then reboot the root domain:

- Add the bus during a delayed reconfiguration
- Add the bus dynamically

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

**Bug ID 16659506:** 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
```

### **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
```

This limitation applies only to the supported SPARC systems prior to the SPARC M7 series servers and SPARC T7 series servers.

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 `::irmpools` and `::irmreqs` MDB macros to determine how interrupts are used. The `::irmpools` macro shows the overall supply of interrupts divided into pools. The `::irmreqs` macro shows which devices are mapped to each pool. For each device, `::irmreqs` 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 `::irmpools` and `::irmreqs` 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 `::irmpools` and `::irmreqs` macros show the following information:

```
# echo "::irmpools" | mdb -k
```

```

ADDR          OWNER  TYPE  SIZE  REQUESTED  RESERVED
00000400016be970 px#0    MSI-X 36    36          36

# echo "00000400016be970::irmreqs" | 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 T5-8 Server: Uptime Data Shows a Value of 0 for Some ldm List Commands

**Bug ID 16068376:** On a SPARC T5-8 server 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.

### **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.

### **`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).

## 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 servers, 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
```

## 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.

## **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.

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

**Bug ID 15701853:** 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.

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

**Bug ID 15668368:** A SPARC T3-1 server 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 server, 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
```

## 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.

## **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.

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

**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.

## **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.

## **The 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
```

## 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.

## 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

## Documentation Issues

This section contains documentation issues and errors that have been found too late to resolve for the Oracle VM Server for SPARC 3.4 release.

### Must Reboot an Active Domain When Using the `ldm set-domain` Command to Change the `boot-policy` Property Value

The [ldm\(1M\)](#) man page does not mention that you must reboot an active domain after you use the `ldm set-domain` command to change the `boot-policy` property value.

The description of the `boot-policy` property has been updated with the following paragraph:

If the domain is active when you change the `boot-policy` value, you must reboot the domain to make the change take effect.

In addition, the first paragraph of the Set Options for Domains section now mentions the `boot-policy` property name:

The `set-domain` subcommand enables you to modify *only* the `boot-policy`, `mac-addr`, `hostid`, `failure-policy`, `extended-mapin-space`, `master`, and `max-cores` properties of each domain. You *cannot* use this command to update resource properties.

### `ldmd(1M)` Man Page Shows an Incorrect SMF Property Name

The [ldmd\(1M\)](#) man page shows the incorrect SMF property name, `ldmd/fj-ppar-dr-policy`. The correct property name is `ldmd/fj_ppar_dr_policy`.

## Resolved Issues

The following enhancement requests and bugs have been fixed for the Oracle VM Server for SPARC 3.4 software release:

---

15736512	ldm add-vsw primary-vsw2 primary "Id already exists" error
15749727	RFE: modifications to hard cap should auto-trigger delayed reconfig
15755622	Path /vdisks/ldg1/disk.img is valid but not accessible on service domain
15778392	Remove workaround that keeps lowest physical strand in primary
15811095	RFE: Improve readability of ls-bindings VSW output
15816108	Ops Center java XMPP client doesn't work on java 1.7.0_07
16417609	Introduce support for "auto" as a value for inter-vnet-link option of a vsw
16762488	Support for Verified Boot
17695061	Need ldm set-vsw command to modify current list of vlans
19073165	Disable support for TLS v1.0 and TLS v1.1
19433898	ovmtdeploy should check for sufficient device space to deploy templates
19926210	Memory add reported disabled after mem range blacklist, primary reboot
19974670	snmp_run_cmd() fails if MIBDIRS is set
20046218	Default vnet linkprop state should be "phys-state"
20457740	Logical Domains Manager should support enabling trusted mode for VNETs
20532270	VHBA: I/O removal should be aware of VSAN devices prior to removing pci bus
20539015	Recovery Mode should support vHBA
20588981	Change number of migration threads to optimize performance
20653017	OVMT Utility support for VLAN tagging
20704553	ldmd should prevent rm-mem if a guest is in transition state

20906346	Deprecate and remove support for the Netra Data Plane SW Suite
21028310	Add range check for MAC address during domain creation
21071542	Monitoring frequency to speed-up migration
21084034	The alt-mac-addr stored in mac_list is not in-sync with vnet
21084239	alt-mac-addr used in failed set-vnet call can't be reused even if not in use
21114622	Logical Domains Manager should not rely on polling to get config status from DIO agent
21116665	Logical Domains Manager provides the failure policy information in the MD for IOR
21116830	ovmtlibrary should have a delete option
21131623	ovmtdeploy should list fully qualified properties including namespace
21179822	Add support for System Domains
21225521	ovmtprop claims to start domain but doesn't try (and probably doesn't need to)
21232477	Live migration fails if vdisk name longer than 31 characters
21232506	ovmtprop prints the FutureWarning: The behavior of this method will change
21283102	ldm list-group displays the same mem or I/O under both /SYS/MB and cpuboard group
21299404	ldm shrink-socket removes extra memory if unaligned mblock
21343710	Remove support for NIU hybrid I/O
21367043	Inconsistent Fujitsu sockets cause fatal error: xalloc(0,4) in affinity_core.c
21369897	ldmpower causes ldmd to SEGV in om_format_pwr_disp()
21385030	Migrate vhba timeout in inactive and bound states

---

21392533	Retire PRI DS between Logical Domains Manager and the SP
21419190	Logical Domains Manager needs to support keystore-backup to persist key values
21419645	Remove thrpool_drain_tasks calls from terminate_mem_mig_loop
21429391	False failure can be returned for an add resource request
21487282	Suppress [current] tag for factory-default in recovery mode
21513120	Logical Domains Manager needs a way to signal GM if an active config is transferred
21527087	ldm set-socket of active domain may cause unstable operation
21532912	Buffer overrun in process_vsan_args of vhba_cmds.c:123
21539695	Domain is left suspended when a suspend request times out
21554591	Migration of multi-mblock on the same LGPG domain on SPARC T7 series servers causes ldmd to crash
21561834	DRM doesn't add vcpus when util=100 em=0 attack=1
21576100	vHBA should be supported by the SNMP MIB
21616749	Page retire request rejected, fma_mem_page_retire routine reported not found
21630889	ldmd core repeatedly in Assertion failed: is_range_contained_memory_list
21635033	ovmcreate fails if service domain has multiple vds
21636259	ldomMIB.so: symbol ldm_priv_init: referenced symbol not found
21653091	Fix vHBA parsed output
21674282	ldm add-vsan fails after replacing the PCI card
21684669	Upgrading to 3.4 causes migration failure for existing bound/active domains

21765169	Vnet with inter-vnet-link=auto will cause corrupted MD on unbind/rebind
21779989	ldm list-hba command shows incorrect aliases of PCI-BOX SLOTS
21780045	ovmcreate generate incorrect ovf file if the locale is not C
21791870	Template memory requirements is failing when it should not
21792094	Dump hashes of guest domain memory to help detect guest memory corruption
21866299	Recovery mode failed with mpgroup dependencies
21876330	OVM Template utilities: combined interface for ovmt admin
21884138	ovmcreate should create templates from domains with virtual disks based on dev
21884198	ovmtlibrary should support templates for disk image other than .img extension
21889318	Memory DR status of an impacted guest may be overwritten
21910643	ovmcreate has a minor formatting error when reporting elapsed time
21911054	teardown_one_vnet_peer_binding() should use CHAIN_ITER_VOL() to remove peers
21926985	ovmtprop get-prop should offer a way to only return the value without key
21936667	Duplicate MAC addr in use message
21946737	Logical Domains Manager ds_netsvc missing pvlan/vlan/vxlan information
21952270	cids_list is not correctly freed
21958105	Add domain dependency support for VHBA
21962779	deleteboard doesn't send proper error message to XSCF in some cases
21969344	set-vnet/set-vsw didn't handle the empty value given to '=+' and '=-'

---

21969537	Add expanded LDC pool support
22018630	ldm set-vnet maxbw accepts out of bound value
22031863	Suspend of a root domain with a large I/O config could time out
22044698	Malformed XML Document failure when the wcore is the only property in ldom_info
22055993	The control domain requires lowest core
22081356	SPARC M7-16 server: ldmd core dump due to fatal error: (1)Getting LDom utilization from the HV
22092431	ldm remove-mem -g fails: not enough free memory present
22106342	Domain incorrectly tagged with an unrecovered suspend error
22150301	"ldm set-vnet max-mac-addr=..." failure message needs to be changed
22156664	Re-enabling custom on an trusted vnet unsupported OS gives confusing outputs
22159550	Need better handling of 'unread HV replies'
22202992	Refactor fatal() et. al. so that fatal function can be marked noreturn
22223877	inter-vnet-link should automatically go off during migration if needed,
22225905	Trusted vnet: pvid/pvlan and custom should be mutually exclusive settings
22233118	ldm list-services wrongly reports VHBA in services list
22258971	Required num. COD permits may be miscalculated on bind of named CMI resources
22263190	Fix leaks and incorrect IOV_VF_FC_NODE_NODE check in xml_v3_resources.c
22289495	mdstore v3.1 causes an argument mismatch in autosave function
22337860	ldm CLI should direct error messages to stderr

22348742	HVctl error code reservations for CMI
22372164	ovmtutils project lost connection to junit tests
22380062	Logical Domains Manager support to online virtual functions on the I/O domain when the root domain comes back up
22381087	Shadow constraints cause bind errors after init-system -f
22385302	getopt() with long options using CLIP is not compatible with GNU getopt()
22448280	Small memory leak in vsw_list_one()
22464339	Assertion failed: argc < create_args, file mig_be_common.c, line 500
22468326	Migration Timeout with Large MD
22468352	Increase DR timeout
22504443	ldm ls-config can crash ldmd when EoUSB is not available
22518200	Attempting a migration causes target ldmd to SEGV in init_pmig_buffers()
22526805	pciv_event_reporting prop is set in an I/O domain
22529488	Reduce suspend time for Board DR Ratio mode
22531007	Dry run of ovmtcreate using the -n option takes same time as an actual run
22541368	Autosave files not updated for vnet/vsw configuration changes
22576504	vHBA: wrong error shown when rm-io removed when vsan is still attached
22587264	Using set-vsw changes inter-vnet-link property from auto to on
22619279	deleteboard with unbind=shutdown doesn't work correctly in ratio mode
22630038	Race condition in IOR context between root domain and I/O domain
22634983	ldmd goes to maintenance during memory DR operations

---

22646805	ovmtdeploy -l gives java.lang.IndexOutOfBoundsException
22655663	ASSERT failure in database_ldom_add_net_client for phys-state value
22661113	ovmtdeploy -M should use the physical address in hex for start-addr
22667042	OVM may core dumped repeatedly if domains are unbound on CPU core deactivation
22670933	ovmtdeploy -d ldg3 -t primary-vsw3 -e net5 template3.ova fails
22672645	ovmtdeploy -d ldg3,ldg4 -N 2 -I pci_3,pci_4 template3.ova fails
22678731	ldmd dumps core on list-io operation post reboot of the primary
22683230	ovmtprop set-prop should give an error instead of Traceback
22683864	ldm set-io should fail if # of alt-mac addrs exceeds max dev configured value
22695268	ovmtutils python agent code get_variables method prints AttributeError
22699679	ovmtlibrary should give an informative error message when no space
22708116	vhba: vhbaTable should use an index into vsan, not a displaystring
22720589	Excessive # warns logged: warning: dr-vio service not registered for guest ldg1
22723793	ovmtcreate does not create the template properly
22728591	../../../../common/src/init.c:336: ABORT Failure to parse initial MD
22751967	iov_allow_domain_suspend() causes PPAR DR regression on M10 platform
22806593	ldmd coredump on iov_pf_macsa_available+0x94
22828719	Temporary fifo file may be not removed from system
22831366	ldmd dumped core on pcie_is_root_domain+0x2c with cmd ls -p ldg0
22844969	"Invalid response" if we use "auto-alt-mac-addrs=0"

22857300	Check for telnet client fails
22862433	Observability for resources allocated to F0 domains in ldm ls-devices o/p
22862736	ldmd core dump in sequence during core retire
22868991	DIOV & DBUS to be treated as static ops if the physical function driver exports any restriction
22875705	ldm set-domain boot-policy should give an error for an active domain
22885443	ldom-db.xml does not show the boot-policy value after a system power-cycle
22896107	"init-system" fails when xml constraint file contains vsan/vhba instances
22911920	Implementation of KEYSTORE_UPDATES_REQUEST/RESPONSE for Guests
22955538	ldmd goes to maintenance when ldmd/nocfg is set to true
22960167	Logical Domains Manager sends RC_READY to I/O domain even if root domain is down
22998250	Harden rprintf() against NULL resp->bufp
22998961	Some guest domains have boot-policy=n/a by default
23020466	Occupied PCIe slots stuck in UNK status on booting up a saved spconfig
23027810	deleteboard with ratio mode may fail if cpu remap is required
23049772	Logical Domains Manager inappropriately applies "diov-capability" property value to bound domains
23060907	Assert failure in debug_chk_mem_align_page() during migration
23075554	Recovery Mode fails to recover I/O domains with available virtual functions
23105178	ldm create-vf command causes ldmd to core dump
23114050	SNMP dumps core in ldomMIB.so:parse_xml_v3_get_ldom_cnt in degraded config

23152267      'Unable to send suspend request' failure during migration

