



Logical Domains (LDoms) 1.1 Release Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 820-4914-10
December 2008, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, Solaris Security Toolkit, JumpStart, OpenBoot, Sun Fire, Netra, Sun Blade, SunSolve, Sun BluePrints, Sun Management Center, Sun Explorer, SPARC, UltraSPARC, and SunVTS are service marks, trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S., or its subsidiaries, and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The Adobe PostScript logo is a trademark of Adobe Systems, Incorporated.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, États-Unis. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuels relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains listés à l'adresse <http://www.sun.com/patents> et un ou les brevets supplémentaires ou les applications de brevet en attente aux États - Unis et dans les autres pays.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, Solaris, JumpStart, OpenBoot, Sun Fire, SunSolve, et SunVTS sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, aux États-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Le logo Adobe PostScript est une marque déposée de Adobe Systems, Incorporated.

Les produits qui font l'objet de ce manuel d'entretien et les informations qu'il contient sont regis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des États-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont regis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DÉCLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES DANS LA LIMITE DE LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

Contents

Logical Domains (LDoms) 1.1 Release Notes 1

What's New in This Release 1

System Requirements 4

 Supported Platforms 4

 Required Software and Patches 5

 Required and Recommended Solaris OS 5

 Required Solaris 10 10/08 Patches 5

 Required Software to Enable LDoms 1.1 Features 6

 Minimum Version of Software Required 6

 Required System Firmware Patches 7

 Location of Logical Domains 1.1 Software 7

 Location of Patches 8

 Location of Documentation 8

Related Software 9

 Additional Recommended Software 9

 Optional Software 9

 Software That Can Be Used With the Logical Domains Manager 10

 System Controller Software That Interacts With Logical Domains Software
 10

Known Issues 11

General Issues	11
Service Processor and System Controller Are Interchangeable Terms	12
Cards Not Supported	12
OpenBoot Firmware Now Supports <code>power-off</code> Command	12
Output of the <code>ldm ls-config</code> Command Changed	12
In Certain Conditions, a Guest Domain's SVM Configuration or Metadevices Can Be Lost	13
Logical Domain Channels (LDCs) and Logical Domains	14
Memory Size Requirements	15
Booting a Large Number of Domains	15
Clearly Shutting Down and Power Cycling a Logical Domains System	16
Memory Size Requested Might Be Different From Memory Allocated	17
Split PCI Regresses in FMA Functionality From Non-Logical Domains Systems	17
Logical Domain Variable Persistence	18
Sun SNMP Management Agent Does Not Support Multiple Domains	19
Do Not Use CPU Power Management If Domains Have Cryptographic Units Bound	19
The <code>sysfwdownload</code> Utility Takes Significantly Longer to Run While LDOMs Is Enabled on Certain Systems	19
Processor Sets and Pools Are Not Compatible With CPU Power Management	19
Bugs Affecting Logical Domains 1.1 Software	20
Logical Domains Manager Can Erroneously Assign an Offline CPU to a Logical Domain	20
Logical Domains Manager Does Not Validate Disk Paths and Network Devices	20
Hang Can Occur With Guest OS in Simultaneous Operations	22
Behavior of the <code>ldm stop-domain</code> Command Can Be Confusing	22
Cannot Set Security Keys With Logical Domains Running	23

Logical Domains Manager Forgets Variable Changes After a Power Cycle	23
Page Retirement Does Not Persist in the Logical Domains Environment	23
Using the <code>server-secure.driver</code> With an NIS Enabled System, Whether or Not LDoms Is Enabled	24
Network Performance Is Substantially Worse in a Logical Domain Guest Than in a Non-LDoms Configuration	25
Logical Domain Time-of-Day Changes Do Not Persist Across a Power Cycle of the Host	26
Harden <code>ds_pri</code> Driver to Protect Against Logical Domains Hanging Indefinitely	26
OpenBoot PROM Variables Cannot be Modified by the <code>eeeprom(1M)</code> Command When the Logical Domains Manager is Running	26
Errors to Buses in a Split-PCI Configuration Might Not Get Logged	26
Emulex-based Fibre Channel Host Adapter Not Supported in Split-PCI Configuration on Sun Fire T1000 Servers	27
Starting and Stopping SunVTS Multiple Times Can Cause Host Console to Become Unusable	27
Virtual Disk Timeouts Do Not Work If Guest or Control Domain Is Halted	27
Logical Domains Manager Does Not Retire Resources On Guest Domain After a Panic and Reboot	28
Guest Domain With Too Many Virtual Networks on the Same Network Using DHCP Can Become Unresponsive	28
Fault Manager Daemon Dumps Core On a Hardened, Single Strand Control Domain	28
Attempting to WAN Boot a Logical Domain Using a Solaris 10 8/07 OS Installation DVD Causes Hang	29
The <code>scadm</code> Command Can Hang Following an SC or SP Reset	29
Domain Services Uses One Lock for Read and Write, Which Can Cause a Logical Domain to Hang	29
Addition of Virtual Disk or Network Devices Under Delayed Reconfiguration Can Fail	29

Simultaneous Net-Installation of Multiple Domains Fails When in a Common Console Group	30
System Can Panic During Reboot If the Virtual Switch Is Configured to Use an Aggregated Network Device	30
Sometimes, the <code>prtdiag(1M)</code> Command Does Not List All CPUs	30
SVM Volumes Built on Slice 2 Fail JumpStart When Used as the Boot Device in a Guest Domain	30
Attempting a Network Interface Operation While Changing a Configuration Using CPR DR Can Cause the Logical Domains Manager to Terminate	31
Deadlock Occurs Rarely With CPU DR Operations	31
If the Solaris 10 5/08 OS is Installed on a Service Domain, Attempting a Net Boot of the Solaris 10 8/07 OS on Any Guest Domain Serviced by It Can Hang the Installation	31
Cryptographic DR Changes Incompatible With Pre-LDoms Firmware	32
ZFS Pool Label Does Not Indicate That a Pool Is Closed Cleanly	32
Sometimes CPU PM Fails to Retrieve PM Policy From the SP	33
Logical Domains Manager Can Take Over 15 Minutes to Shut Down a Logical Domain	33
Rarely, Control Domain Can Panic While Dynamically Removing a Virtual Network Interface From a Domain	33
Virtual Network Using Hybrid I/O Might Not Tag or Untag Packets When the Port VLAN ID Is Set	34
Solaris 10/08 OS Installation Hangs With ZFS Boot on Certain Servers With 1-GB Memory	34
Do Not Switch to Performance Mode in CPU Power Management Unless All Domains Are Up and Running the Solaris OS	34
With Elara Copper Card, the Service Domain Hangs on Reboot	34
Sometimes, Executing the <code>uadmin 1 0</code> Command From an LDoms System Does Not Return the System to the OK Prompt	35
Unrecoverable Hardware Error Seen on Guest With 2 Hybrid Virtual Networks When the Service Domain Stops	35
Logical Domains Manager Displays Migrated Domains in Transition States When They Are Already Booted	36

Logical Domains Manager Does Not Start If the Machine Is Not Networked and an NIS Client Is Running	36
Newly Configured Virtual Network Fails to Establish a Connection With the Virtual Switch	36
Do Not Migrate a Guest Domain That Is at the <code>kmdb</code> Prompt	37
Cannot Export a ZFS Volume as a Single-Slice Virtual Disk From Service Domain Running Solaris 10 5/08 OS, or Earlier, to Guest Domain Running Solaris 10 10/08 OS	37
Migration Can Fail to Bind Memory Even If the Target Has Enough Available	37
Migration Does Not Fail If a <code>vdsdev</code> on the Target Has a Different Backend	38
After Node Reconfiguration, Logical Domains Manager Might Fail to Start on SPARC Enterprise T5440 Servers	38
Constraint Database Is Not Synchronized to Saved Configuration	39
Cryptic Message to Indicate Incoming Migration Is Disabled on the Target	39
Two Faulty CPUs That Include CPU 0 Can Cause Power Management Utilization Issue After Rebooting	39
Domain Can Lose Virtual CPUs During a Migration If Another Domain Is Rebooting	40
Panic Occurs When More Than Two Hybrid I/O–Capable Virtual Networks Are Activated on a Guest Domain	40
Migration Does Not Clean Up on Target If Virtual Network MAC Address Clashes With Existing Domain	41
Explicit Console Group and Port Bindings Are Not Migrated	41
Migration Dry Run Does Not Detect Inadequate Memory	41
In Certain Situations, a Domain Migration Can Fail After the Control Domain Is Rebooted	41
Pseudonyms for PCI Busses on Sun SPARC Enterprise T5440 Systems Are Not Correct	42
Cancelling Domain Migration With Virtual Networks Using Multiple Virtual Switches Might Cause Domain Reboot	42

Possibility That New Domain on Source System Could Be Assigned Same MAC Address as Domain That Was Successfully Migrated	43
Documentation Errata	44
VIO DR Operations Ignore the Force (-f) Option	44
The <code>ldm rm-reconf</code> Command No Longer Works	44
The <code>vntsd(1M)</code> Man Page Description for <code>listen-addr</code> Is Incomplete	44
Error in Descriptions of <code>default-vlan-id</code> , <code>pvid</code> , and <code>vid</code> in the <code>ldm(1M)</code> Man Page	45
The <code>status</code> Output Option Was Omitted for the <code>list</code> Subcommand in the <code>ldm(1M)</code> Man Page	45
Resolved Issues	46
LDoms 1.1 RFEs and Bugs Fixed in Solaris 10 10/08 OS	46
RFEs and Bugs Fixed for LDoms 1.1 Software	48

Logical Domains (LDoms) 1.1 Release Notes

These release notes contain changes for this release, a list of supported platforms, a matrix of required software and patches, and other pertinent information, including bugs that affect Logical Domains 1.1 software.

What's New in This Release

The major changes for this release of Logical Domains 1.1 software are as follows:

- Extensible Markup Language (XML) version 3 (v3) public interface for monitoring and controlling logical domains provides:
 - Access using the Extensible Messaging and Presence Protocol (XMPP)
 - Local raw access with no authentication
 - Logical Domains Manager interface and Logical Domains Manager events(Refer to Chapter 10 of the *Logical Domains (LDoms) 1.1 Administration Guide* for details about using this feature.)

- Guest Domain Migration provides the capability to migrate a guest logical domain from one server to another compatible server. If active, the domain on the source server is suspended, and its configuration and run-time state are transferred to another server, where the domain is recreated and resumed. Inactive and bound domains can also be migrated, and this happens almost instantaneously as only the domain's configuration needs to be transferred and recreated. Refer to Chapter 8, "Migrating Logical Domains" in the *Logical Domains (LDoms) 1.1 Administration Guide* for more information.
- Virtual Input/Output (I/O) Dynamic Reconfiguration (DR) provides the ability to add and remove virtual I/O services and devices without rebooting.
- Network
 - Network Interface Unit (NIU) Hybrid I/O provides support for a virtualized I/O path, a hybrid I/O path, and better performance and scalability. Refer to "Using NIU Hybrid I/O" in the *Logical Domains (LDoms) 1.1 Administration Guide* for more information.
 - VLAN Support provides the capability to configure and use virtual local area networks (VLANs) in logical domains. Refer to "Using VLAN Tagging With Logical Domains Software" in the *Logical Domains (LDoms) 1.1 Administration Guide* for how to use this feature.

Note – LDoms 1.1 software and Solaris™ 10 10/08 OS are required to use this feature. Tagged VLANs are not supported in any of the previous releases for LDoms networking components.

- Storage
 - Virtual Disk Failover adds support for disk multipathing which enables the virtual disk device in a guest domain to be serviced by multiple virtual disk servers. Refer to "Configuring Virtual Disk Multipathing" in the *Logical Domains (LDoms) 1.1 Administration Guide* for information about how to set up this feature.
 - Single-Slice Disk Enhancement allows installing the Solaris OS on a single-slice disk. Single-slice disks are now visible with the `format(1M)` command.
- New and Changed CLI `ldm` Subcommands — refer to the `ldm(1M)` man page or the *Logical Domains (LDoms) Manager 1.1 Man Page Guide* for more details.
 - Use the new `list` subcommand output (`-o format`) options to limit the output format to one or more of the following subsets: `console`, `cpu`, `crypto`, `disk`, `domain`, `memory`, `network`, `physio`, `serial`, and `status`.
 - Use new arguments to the `add-vnet` and `set-vnet` subcommands to specify details for hybrid I/O (`mode=hybrid`).
 - Use new arguments to the `add-vsw`, `set-vsw`, `add-vnet`, and `set-vnet` subcommands to specify VLAN tagging (port VLAN ID and VLAN ID).

- Use the new `mpgroup=` argument to the `add-vdsdev` and `set-vdsdev` subcommands to define a multipath group name for several virtual disk server devices (`vdsdev`). In case a virtual disk cannot communicate with a virtual disk server device, a failover is initiated to one of the other virtual disk server devices in the multipath group.
- Use the `ldm cancel-operations reconf` command instead of the deprecated `ldm remove-reconf` or `cancel-reconf` commands to cancel a delayed reconfiguration operation. You can still use the `remove-reconf` or `cancel-reconf` subcommands as an alias, but the `rm-reconf` subcommand does not work.
- Use the `ldm cancel-operations migration` command to cancel a domain migration operation.
- Use the `ldm migrate-domain` command to migrate a logical domain from one machine to another.
- Output of `ldm ls-config` has changed so that the annotations more accurately reflect when a configuration saved to the service processor (SP) matches the currently running configuration. See [“Output of the `ldm ls-config` Command Changed” on page 12](#) for more information.
- CPU Power Management (PM) implements power saving by disabling each Sun UltraSPARC® T2- and T2 Plus-based Server core that has all of its CPU strands idle. Refer to *“Using CPU Power Management With LDomS 1.1 Software”* in the *Logical Domains (LDoms) 1.1 Administration Guide*. To use CPU power management, your system must have the firmware that supports power management.
- Bug Fixes

System Requirements

This section contains system requirements for running LDom software.

Supported Platforms

Logical Domains (LDoms) Manager 1.1 software is supported on the following platforms:

TABLE 1 Supported Platforms

Name	Reference
Sun UltraSPARC T2 Plus–based Servers:	
Sun SPARC® Enterprise T5140 and T5240 Servers	<i>Sun SPARC Enterprise T5140 and T5240 Servers Administration Guide</i>
Sun SPARC Enterprise T5440 Server	<i>Sun SPARC Enterprise T5440 Server Administration Guide</i>
Sun Blade™ T6340 Server Module	<i>Sun Blade T6340 Server Module Product Notes</i>
Netra™ T5440 Server	<i>Sun Netra T5440 Server Product Notes</i>
Sun UltraSPARC T2–based Servers:	
Sun SPARC Enterprise T5120 and T5220 Servers	<i>Sun SPARC Enterprise T5120 and T5220 Servers Administration Guide</i>
Sun Blade T6320 Server Module	<i>Sun Blade T6320 Server Module Product Notes</i>
Netra CP3260 Blade	<i>Netra CP3260 Board Product Notes</i>
Netra T5220 Server	<i>Sun Netra T5220 Server Product Notes</i>
Sun UltraSPARC T1–based Servers:	
Sun Fire™ or SPARC Enterprise T1000 Server	<i>Sun Fire or SPARC Enterprise T1000 Server Administration Guide</i>
Sun Fire or SPARC Enterprise T2000 Server	<i>Sun Fire or SPARC Enterprise T2000 Server Administration Guide</i>
Netra T2000 Server	<i>Netra T2000 Server Administration Guide</i>
Netra CP3060 Blade	<i>Netra CP3060 Board Product Notes</i>
Sun Blade T6300 Server Module	<i>Sun Blade T6300 Server Module Administration Guide</i>

Required Software and Patches

This section lists the required software and patches for use with Logical Domains 1.1 software.

Required and Recommended Solaris OS

To use all the features of LDomS 1.1 software, the operating system on all domains should be at least equivalent to the Solaris 10 10/08 OS. This can be either a fresh or upgrade installation of the following:

- Solaris 10 10/08 OS
- Solaris 10 5/08 OS with Patch ID 137137-09
- Solaris 10 8/07 OS with Patch ID 137137-09
- Solaris 10 11/06 OS with Patch ID 137137-09

Required Solaris 10 10/08 Patches

Following are the required Solaris 10 10/08 patches for use with Logical Domains 1.1 software. An X marks whether a patch must be installed on that specific type of domain, but the patches can be applied to all domains.

TABLE 2 Required Solaris 10 10/08 Patches and Domains Needing Patch

Patch ID	Control Domain	Service-I/O Domain	Guest Domain
139458-01 (aggr driver)	X	X	
139502-01 (picl plugin)	X		
139508-01 (niumx driver)	X	X	X
139562-02 (multiple LDomS drivers)	X	X	X
139570-02 (nxge driver)	X	X	X

Required Software to Enable LDom 1.1 Features

Following is a matrix of required software to enable all the Logical Domains 1.1 features.

TABLE 3 Required Software to Enable Logical Domains 1.1 Features

Supported Servers	System Firmware	Solaris OS
Sun UltraSPARC T2 Plus-based servers	7.2.x	One of the configurations in "Required and Recommended Solaris OS" on page 5
Sun UltraSPARC T2-based servers	7.2.x	One of the configurations "Required and Recommended Solaris OS" on page 5
Sun UltraSPARC T1-based servers	6.7.x	One of the configurations "Required and Recommended Solaris OS" on page 5

Minimum Version of Software Required

It is possible to run the Logical Domains 1.1 software along with previous revisions of the other software components. For example, you could have differing versions of the Solaris OS on the various domains in a machine. It is recommended to have all domains running Solaris 10 10/08 OS plus the patches listed in [TABLE 2](#). However, an alternate upgrade strategy could be to upgrade the control and service domains to Solaris 10 10/08 OS plus the patches listed in [TABLE 2](#) and to continue running the guest domains at the existing patch level.

Following is a matrix of the minimum versions of software required. The LDom 1.1 package, `SUNWldm`, can be applied to a system running at least the following versions of software. The minimum software versions are platform specific and depend on the requirements of the CPU in the machine. The minimum Solaris OS version for a given CPU type applies to all domain types (control, service, I/O, and guest).

TABLE 4 Minimum Versions of Software

Supported Servers	System Firmware	Solaris OS
Sun UltraSPARC T2 Plus-based servers	7.1.x	Solaris 10 8/07 plus patch ID 127111-08 at a minimum
Sun UltraSPARC T2-based servers	7.0.x	Solaris 10 8/07
Sun UltraSPARC T1-based servers	6.5.x	Solaris 10 11/06 plus patch IDs 124921-02, 125043-01, and KU 118833-36 at a minimum

Required System Firmware Patches

Following are the required system firmware patches at a minimum for use with Logical Domains 1.1 software on supported servers.

TABLE 5 Required System Firmware Patches

Patches	Supported Servers
139434-01	Sun Fire and SPARC Enterprise T2000 Servers
139435-01	Sun Fire and SPARC Enterprise T1000 Servers
139436-01	Netra T2000 Server
139437-01	Netra CP3060 Blade
139438-01	Sun Blade T6300 Server Module
139439-01	Sun SPARC Enterprise T5120 and T5220 Servers
139440-01	Sun Blade T6320 Server Module
139441-01	Netra CP3260 Blade
139442-01	Netra T5220 Server
139444-01	Sun SPARC Enterprise T5140 and T5240 Servers
139445-01	Netra T5440 Server
139446-01	Sun SPARC Enterprise T5440 Server
139448-01	Sun Blade T6340 Server Module

Note – The -01 versions of the system firmware Patch IDs do not support power management.

Location of Logical Domains 1.1 Software

You can find the LDOMs 1.1 software to download from the following web site:

<http://www.sun.com/ldoms>

The `LDoms_Manager-1_1.zip` file that you download contains the following:

- Logical Domains Manager 1.1 software (`SUNWldm.v`).
- The `ldm(1M)` man page is included in the `SUNWldm.v` package and gets installed when the package is installed.
- Installation script for Logical Domains Manager 1.1 software and the Solaris Security Toolkit (`install-ldm`)

- Solaris Security Toolkit (SUNWjass)
- Logical Domains Management Information Base (SUNWldmib.v)
- Libvirt source file (Libvirt-source)
- Libvirt for Logical Domains (SUNWldlibvirt.v)
- Install package for Libvirt for Logical Domains (SUNWldvirtinst.v)

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

```

LDoms_Manager-1_1/
  Install/
    install-ldm
  Legal/
    LDoms_1.0.1_libvirt_entitlement(20071220).txt
    Ldoms_1.1_Entitlement.txt
    Ldoms_1.1_SLA_Entitlement.txt
    Ldoms_MIB_1.0.1_Entitlement.txt
    Ldoms_MIB_1.0.1_SLA_Entitlement.txt
    LGPLDisclaimer.txt
    THIRDPARTYREADME(20071220).txt
  Product/
    SUNWjass
    SUNWldm.v
    SUNWldvirtinst.v
    Libvirt-source
    SUNWldlibvirt.v
    SUNWldmib.v
  README

```

Location of Patches

You can find the required Solaris OS and system firmware patches at the SunSolveSM site:

<http://sunsolve.sun.com>

Location of Documentation

The *Logical Domains (LDoms) 1.1 Administration Guide* and *Logical Domains (LDoms) 1.1 Release Notes* can be obtained from:

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

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

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

The *Beginners Guide to LDom: Understanding and Deploying Logical Domains* can be used to get a general overview of Logical Domains software. However, the details of the guide specifically apply to the LDom 1.0 software release and are now out of date for LDom 1.1 software. The guide can be found at the Sun BluePrints™ site.

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

Related Software

This section describes software that is related to LDom software.

Additional Recommended Software

Solaris Security Toolkit 4.2 software – This software can help you secure the Solaris OS in the control domain and other domains. Refer to the *Solaris Security Toolkit 4.2 Administration Guide* and *Solaris Security Toolkit 4.2 Reference Manual* for more information.

Optional Software

- **Logical Domains (LDom) Management Information Base (MIB) software** – This software can help you enable third party applications to perform remote monitoring and a few control operations. Refer to the *Logical Domains (LDom) MIB 1.0.1 Administration Guide* and *Release Notes* for more information.
- **Libvirt for LDom software** – This software provides virtual library (`libvirt`) interfaces for Logical Domains (LDom) software so that virtualization customers can have consistent interfaces. The `libvirt` library (version 0.3.2) included in this software interacts with the Logical Domains Manager software running on Solaris 10 Operating System (OS) to support Logical Domains virtualization technology. Refer to the *Libvirt for LDom 1.0.1 Administration Guide* and *Release Notes* for more information.

Note – LDom MIB software and Libvirt for LDom software works with LDom 1.0.1 software at a minimum.

Software That Can Be Used With the Logical Domains Manager

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

- **SunVTS** functionality is available in the control domain and guest domains on certain LDom software releases and certain platforms. SunVTS™ is Sun's Validation Test Suite, which provides a comprehensive diagnostic tool that tests and validates Sun hardware by verifying the connectivity and proper functioning of most hardware controllers and devices on Sun servers. For more information about SunVTS, refer to the *SunVTS User's Guide* for your version of SunVTS.
- **Sun Management Center 4.0 Add-On Software** can be used only on the control domain with the Logical Domains Manager software enabled. Sun™ Management Center is an open, extensible system monitoring and management solution that uses Java™ and a variant of the Simple Network Management Protocol (SNMP) to provide integrated and comprehensive enterprise-wide management of Sun products and their subsystem, component, and peripheral devices. Support for hardware monitoring within the Sun Management Center environment is achieved through the use of appropriate hardware server module add-on software, which presents hardware configuration and fault reporting information to the Sun Management Center management server and console. Refer to the *Sun Management Center 4.0 Add-On Software Release Notes: For Sun Fire, SunBlade, Netra, and SunUltra Systems* for more information about using Sun Management Center 4.0 on the supported servers.
- **Sun Explorer Data Collector** can be used with the Logical Domains Manager software enabled on the control domain. Sun™ Explorer is a diagnostic data collection tool. The tool comprises shell scripts and a few binary executables. Refer to the *Sun Explorer User's Guide* for more information about using the Sun Explorer Data Collector.
- **Solaris Cluster** software can be used only on an I/O domain in Logical Domains software releases up through LDom 1.0.2. In LDom 1.0.3 and 1.1 software, Solaris™ Cluster software can be used in a guest domain with some restrictions. Refer to Solaris Cluster documentation for more information about any restrictions and about the Solaris Cluster software in general.

System Controller Software That Interacts With Logical Domains Software

The following system controller (SC) software interacts with the Logical Domains 1.1 software:

- **Sun Integrated Lights Out Manager (ILOM) 3.0** firmware is the system management firmware you can use to monitor, manage, and configure Sun UltraSPARC T2-based server platforms. ILOM is preinstalled on these platforms and can be used on the control domain on LDoms-supported servers with the Logical Domains Manager 1.1 software enabled. Refer to the *Sun Integrated Lights Out Manager 3.0 User's Guide* for features and tasks that are common to Sun rackmounted servers or blade servers that support ILOM. Other user documents present ILOM features and tasks that are specific to the server platform you are using. You can find the ILOM platform-specific information within the documentation set that accompanies your system.
- **Advanced Lights Out Manager (ALOM) Chip Multithreading (CMT) Version 1.3** software can be used on the control domain on UltraSPARC® T1-based servers with the Logical Domains Manager 1.0.1 software enabled. Refer to “Using LDoms With ALOM CMT” in the *Logical Domains (LDoms) 1.1 Administration Guide*. The ALOM system controller enables you to remotely manage and administer your supported CMT servers. ALOM enables you to monitor and control your server either over a network or by using a dedicated serial port for connection to a terminal or terminal server. ALOM provides a command-line interface that you can use to remotely administer geographically distributed or physically inaccessible machines. For more information about using ALOM CMT Version 1.3 software, refer to the *Advanced Lights Out Management (ALOM) CMT v1.3 Guide*.
- **Netra Data Plane Software Suite** is a complete board software package solution. The software provides an optimized rapid development and runtime environment on top of multistrand partitioning firmware for Sun CMT platforms. The Logical Domains Manager contains some `ldm` subcommands (`add-vdpcs`, `rm-vdpcs`, `add-vdpcc`, and `rm-vdpcc`) for use with this software. Refer to the *Netra Data Plane Software Suite* documentation for more information about this software.

Known Issues

This section contains general issues and specific bugs concerning the Logical Domains 1.1 software.

General Issues

This section describes general known issues about this release of LDoms software that are broader than a specific bug number. Workarounds are provided where available.

Service Processor and System Controller Are Interchangeable Terms

For discussions in Logical Domains documentation, the terms service processor (SP) and system controller (SC) are interchangeable.

Cards Not Supported

The following cards are *not* supported for this LDomS 1.1 software release:

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



Caution – If these unsupported configurations are used with LDomS 1.1, stop and unbind all logical domains before the control domain is rebooted. Failure to do so can result in a system crash causing the loss of all the logical domains that are active in the system.

OpenBoot Firmware Now Supports `power-off` Command

The OpenBoot™ firmware now supports the `power-off` command. The command powers off the system if only the control domain is active. The OpenBoot `power-off` command behaves on the control domain exactly the way the Solaris OS `halt` command behaves. Refer to Table 9-1 in Chapter 9 of the *Logical Domains (LDoms) 1.1 Administration Guide* for a specific description of the behavior of the `halt` command.

Output of the `ldm ls-config` Command Changed

The output of the `ldm ls-config` command now more accurately reflects when a saved configuration matches the currently running configuration.

- Previously, the configuration that was last booted (that is, on the previous power on) was always listed as `[current]`. Now, the last booted configuration is listed as `[current]` only until you initiate a reconfiguration. After the reconfiguration, the annotation changes to `[next poweron]`.
- Previously, the result of an `ldm add-config` or `set-config` command was that the specified configuration was labeled as `[next]`. Now, such a configuration is listed as `[current]`, because it does match the currently running configuration.
- The `[next]` annotation has changed to `[next poweron]`.

In Certain Conditions, a Guest Domain's SVM Configuration or Metadevices Can Be Lost

If a service domain is running a version of Solaris 10 OS prior to Solaris 10 10/08 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 Solaris 10 10/08 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, this can cause the Solaris™ Volume Manager (SVM) to be unable to find its configuration or to access its metadevices.

Workaround: After upgrading a service domain to Solaris 10 10/08, if a guest domain is unable to find its SVM configuration or its metadevices, execute the following procedure.

▼ Find a Guest Domain's SVM Configuration or Metadevices

1. Boot the guest domain.

2. Disable the `devid` feature of SVM by adding the following lines to the `/kernel/dr/md.conf` file:

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

3. Reboot the guest domain.

After the domain has booted, the SVM configuration and metadevices should be available.

4. Check the SVM configuration and ensure it is correct.

5. Re-enable the SVM `devid` feature by removing the two lines added in Step 2 from the `/kernel/drv/md.conf` file.

6. Reboot the guest domain.

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

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

These messages are normal and do not report any problems.

Logical Domain Channels (LDCs) and Logical Domains

There is a limit to the number of LDCs available in any logical domain. For Sun UltraSPARC T1-based platforms, that limit is 256; for all other platforms, the limit is 512. Practically speaking, this only becomes an issue on the control domain, because the control domain has at least part, if not all, of the I/O subsystem allocated to it, and because of the potentially large number of LDCs created for both virtual I/O data communications and the Logical Domains Manager control of the other logical domains.

Note – The examples in this section are what happens on Sun UltraSPARC T1-based platforms. However, the behavior is the same if you go over the limit on other supported platforms.

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

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

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

1. The control domain allocates 12 LDCs for various communication purposes with the hypervisor, Fault Management Architecture (FMA), and the system controller (SC), independent of the number of other logical domains configured.
2. The control domain allocates one LDC to every logical domain, including itself, for control traffic.
3. Each virtual I/O service on the control domain consumes one LDC for every connected client of that service.

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

- Virtual network
- Virtual disk
- Virtual console

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

$12(1) + 9(2) + 8 \times 3(3) = 45$ LDCs in total.

Now consider the case where there are 32 domains instead of 8, and each domain includes 3 virtual disks, 3 virtual networks, and a virtual console. Now the equation becomes:

$12 + 33 + 32 \times 7 = 269$ LDCs in total.

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

Memory Size Requirements

Logical Domains software does not impose a memory size limitation when creating a domain. The memory size requirement is a characteristic of the guest operating system. Some Logical Domains functionality might not work if the amount of memory present is less than the recommended size. For recommended and minimum size memory requirements, refer to the installation guide for the operating system you are using. Refer to “System Requirements and Recommendations” in the *Solaris 10 Installation Guide: Planning for Installation and Upgrade*.

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

Booting a Large Number of Domains

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

- Up to 128 on Sun UltraSPARC T2 Plus-based servers
- Up to 64 on Sun UltraSPARC T2-based servers
- Up to 32 on Sun UltraSPARC T1-based servers

If unallocated virtual CPUs are available, assign them to the service domain to help process the virtual I/O requests. Allocate 4 to 8 virtual CPUs to the service domain when creating more than 32 domains.

In cases where maximum domain configurations have only a single CPU in the service domain, do not put unnecessary stress on the single CPU when configuring and using the domain.

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

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

- 128 gigabytes of memory for Sun UltraSPARC T2 Plus-based servers
- 64 gigabytes of memory for Sun UltraSPARC T2-based servers
- 32 gigabytes of memory for Sun UltraSPARC T1-based servers

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

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

Cleanly Shutting Down and Power Cycling a Logical Domains System

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

▼ Power Off a System With Multiple Active Domains

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

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

▼ Power Cycle the System

1. **Shut down and unbind all the non-I/O domains.**
2. **Shut down and unbind any active I/O domains.**

3. Reboot the primary domain.

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

Memory Size Requested Might Be Different From Memory Allocated

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

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

Split PCI Regresses in FMA Functionality From Non-Logical Domains Systems

Currently, Fault Management Architecture (FMA) diagnosis of I/O devices in a Logical Domains environment might not work correctly. The problems are:

- Input/output (I/O) device faults diagnosed in a non-control domain are not logged on the control domain. These faults are only visible in the logical domain that owns the I/O device.
- I/O device faults diagnosed in a non-control domain are not forwarded to the system controller. As a result, these faults are not logged on the SC and there are no fault actions on the SC, such as lighting of light-emitting diodes (LEDs) or updating the dynamic field-replaceable unit identifiers (DFRUIDs).
- Errors associated with a root complex that is not owned by the control domain are not diagnosed properly. These errors can cause faults to be generated against the diagnosis engine (DE) itself.

Logical Domain Variable Persistence

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

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

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

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

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

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

In Logical Domains 1.1 software, there are a few cases where variable updates do not persist as expected:

- With domaining enabled (the default in all cases except the UltraSPARC T1000 and T2000 systems running in `factory-default` configuration), all methods of updating a variable (OpenBoot firmware, `eeeprom` command, `ldm` subcommand) persist across reboots of that domain, but not across a power cycle of the system, unless a subsequent logical domain configuration is saved to the SC. In addition, in the control domain, updates made using OpenBoot firmware persist across a power cycle of the system; that is, even without subsequently saving a new logical domain configuration to the SC.
- When domaining is not enabled, variable updates specified through the `eeeprom(1M)` command persist across a reboot of the primary domain into the same `factory-default` configuration, but do not persist into a configuration saved to the SC. Conversely, in this scenario, variable updates specified using the Logical Domains Manager do not persist across reboots, but are reflected in a configuration saved to the SC.

So, when domaining is not enabled, if you want a variable update to persist across a reboot into the same factory-default configuration, use the eeprom command. If you want it saved as part of a new logical domains configuration saved to the SC, use the appropriate Logical Domains Manager command.

- In all cases, when reverting to the factory-default configuration from a configuration generated by the Logical Domains Manager, all LDom variables start with their default values.

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

Sun SNMP Management Agent Does Not Support Multiple Domains

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

Do Not Use CPU Power Management If Domains Have Cryptographic Units Bound

Do not use the CPU power management feature in Integrated Lights-Out Management (ILOM) if your domains are to have cryptographic units bound.

The `sysfwdownload` Utility Takes Significantly Longer to Run While LDom Is Enabled on Certain Systems

The `sysfwdownload` utility takes significantly longer to run from within a Logical Domains environment on systems based on UltraSPARC T1 processors. This happens if you use the `sysfwdownload` utility while the LDom software is enabled.

Workaround: Boot to the `factory-default` configuration with the LDom software disabled before using the utility.

Processor Sets and Pools Are Not Compatible With CPU Power Management

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

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

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

Bugs Affecting Logical Domains 1.1 Software

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

Logical Domains Manager Can Erroneously Assign an Offline CPU to a Logical Domain

Bug ID 6431107: When the Fault Management Architecture (FMA) places a CPU offline, it records that information, so that when the machine is rebooted the CPU remains offline. The offline designation persists in a non-Logical Domains environment.

However, in a Logical Domains environment, this persistence is not always maintained for CPUs in guest domains. The Logical Domains Manager does not currently record data on fault events sent to it. This means that a CPU in a guest domain that has been marked as faulty, or one that was not allocated to a logical domain at the time the fault event is replayed, can subsequently be allocated to another logical domain with the result that it is put back online.

Logical Domains Manager Does Not Validate Disk Paths and Network Devices

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

Disk Paths

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

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

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

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

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

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

▼ Recover From an Errant Virtual Disk Service Device or Volume

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

Network Devices

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

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

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

▼ Recover From a Disk Device Being Used by Other Software

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

Hang Can Occur With Guest OS in Simultaneous Operations

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

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

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

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

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

If the Solaris OS is halted on the domain; for example, by using the `halt(1M)` command; and the domain is at the prompt `"r)eboot, o)k prompt, h)alt?,"` the `ldm stop-domain` command fails with the following error message:

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

Workaround: Force a stop by using the `ldm stop-domain` command with the `-f` option.

Recovery: If you restart the domain from the `kldb` prompt, the stop notification is handled, and the domain does stop.

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

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

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

Cannot Set Security Keys With Logical Domains Running

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

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

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

Logical Domains Manager Forgets Variable Changes After a Power Cycle

Bug ID 6590259: This issue is summarized in [“Logical Domain Variable Persistence” on page 18](#).

Page Retirement Does Not Persist in the Logical Domains Environment

Bug ID 6531058: When a memory page of a guest domain is diagnosed as faulty, the Logical Domains Manager retires the page in the logical domain. If the logical domain is stopped and restarted again, the page is no longer in a retired state.

The `fmadm faulty -a` command shows whether the page from either the control or guest domain is faulty, but the page is not actually retired. This means the faulty page can continue to generate memory errors.

Workaround: Use the following command in the control domain to restart the fault manager daemon, `fmd(1M)`:

```
primary# svcadm restart fmd
```

Using the `server-secure.driver` With an NIS Enabled System, Whether or Not LDoms Is Enabled

Bug ID 6533696: On a system configured to use the Network Information Services (NIS) or NIS+ name service, if the Solaris Security Toolkit software is applied with the `server-secure.driver`, NIS or NIS+ fails to contact external servers. A symptom of this problem is that the `ypwhich(1)` command, which returns the name of the NIS or NIS+ server or map master, fails with a message similar to the following:

```
Domain atlas some.atlas.name.com not bound on nis-server-1.c
```

The recommended Solaris Security Toolkit driver to use with the Logical Domains Manager is `ldm_control-secure.driver`, and NIS and NIS+ work with this recommended driver.

If you are using NIS as your name server, you cannot use the Solaris Security Toolkit profile `server-secure.driver`, because you may encounter Solaris OS Bug ID 6557663, IP Filter causes panic when using `ipnat.conf`. However, the default Solaris Security Toolkit driver, `ldm_control-secure.driver`, is compatible with NIS.

▼ Recover by Resetting Your System

1. Log in to the system console from the system controller, and if necessary, switch to the ALOM mode by typing:

```
# #.
```

2. Power off the system by typing the following command in ALOM mode:

```
sc> poweroff
```

3. Power on the system.

```
sc> poweron
```


4. Switch to the console mode at the `ok` prompt:

```
sc> console
```

5. Power on the system.

```
ok boot -s
```

6. Edit the file `/etc/shadow`, and change the first line of the shadow file that has the root entry to:

```
root::6445::::::
```

7. Log in to the system and do one of the following:

- Add file `/etc/ipf/ipnat.conf`
- Undo the Solaris Security Toolkit, and apply another driver.

```
# /opt/SUNWjass/bin/jass-execute -ui  
# /opt/SUNWjass/bin/jass-execute -a ldm_control-secure.driver
```

Network Performance Is Substantially Worse in a Logical Domain Guest Than in a Non-LDoms Configuration

Bug ID 6486234: The virtual networking infrastructure adds additional overhead to communications from a logical domain. All packets are sent through a virtual network device, which, in turn, passes the packets to the virtual switch. The virtual switch then sends the packets out through the physical device. The lower performance is seen due to the inherent overheads of the stack.

Workaround: Do one of the following depending on your server:

- On Sun UltraSPARC T1-based servers, such as the Sun Fire T1000 and T2000, and Sun UltraSPARC T2+ based servers such as the Sun SPARC Enterprise T5140 and T5240, assign a physical network card to the logical domain using a split-PCI configuration. For more information, refer to “Configuring Split PCI Express Bus to Use Multiple Logical Domains” in the *Logical Domains (LDoms) 1.1 Administration Guide*.
- On Sun Ultra SPARC T2-based servers, such as the Sun SPARC Enterprise T5120 and T5220 servers, assign a Network Interface Unit (NIU) to the logical domain.

Logical Domain Time-of-Day Changes Do Not Persist Across a Power Cycle of the Host

Bug ID 6590259: If the time or date on a logical domain is modified, for example using the `ntpdate` command, the change persists across reboots of the domain but not across a power cycle of the host.

Workaround: For time changes to persist, save the configuration with the time change to the SC and boot from that configuration.

Harden `ds_pri` Driver to Protect Against Logical Domains Hanging Indefinitely

Bug ID 6538932: This requested fix is to attempt to prevent `ds_pri` hangs caused by other bugs. Currently, there are no outstanding bugs that are known to cause `ds_pri` hangs.

Workaround: If a logical domain should hang, stop and restart the affected domain.

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

Bug ID 6540368: This issue is summarized in [“Logical Domain Variable Persistence” on page 18](#) and affects only the control domain.

Errors to Buses in a Split-PCI Configuration Might Not Get Logged

Bug ID 6542295: During operations in a split-PCI configuration, if a bus is unassigned to a domain or is assigned to a domain but not running the Solaris OS, any error in that bus or any other bus might not get logged. Consider the following example:

In a split-PCI configuration, Bus A is not assigned to any domain, and Bus B is assigned to the primary domain. In this case, any error that occurs on Bus B might not be logged. (The situation occurs only during a short time period.) The problem resolves when the unassigned Bus A is assigned to a domain and is running the Solaris OS, but by then some error messages might be lost.

Workaround: When using a split-PCI configuration, quickly verify that all buses are assigned to domains and are running the Solaris OS.

Emulex-based Fibre Channel Host Adapter Not Supported in Split-PCI Configuration on Sun Fire T1000 Servers

Bug ID 6544004: The following message appears at the `ok` prompt if an attempt is made to boot a guest domain that contains an Emulex-based Fibre Channel host adapter (Sun Part Number 375-3397):

```
ok> FATAL:system is not bootable, boot command is disabled
```

Workaround: Do not use this adapter in a split-PCI configuration on Sun Fire T1000 servers.

Starting and Stopping SunVTS Multiple Times Can Cause Host Console to Become Unusable

Bug ID 6549382: If SunVTS™ is started and stopped multiple times, it is possible that switching from the SC console to the host console, using the `console SC` command can result in either of the following messages being repeatedly emitted on the console:

```
Enter #. to return to ALOM.
```

```
Warning: Console connection forced into read-only mode
```

Recovery: Reset the SC using the `resetsc` command.

Virtual Disk Timeouts Do Not Work If Guest or Control Domain Is Halted

Bug ID 6589660: Virtual disk timeouts do not work if either the guest or control domain using the disk is halted; for example, if the domain is taken into the kernel debugger (`kmdb`) or taken into the OpenBoot PROM with the `send break`.

Workaround: None.

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

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

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

```
primary# svcadm restart fmfd
```

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

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

Workaround: Avoid such configurations.

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

Fault Manager Daemon Dumps Core On a Hardened, Single Strand Control Domain

Bug ID 6604253: If you run the Solaris 10 11/06 OS and you harden drivers on the primary domain that is configured with only one strand, rebooting the primary domain or restarting the fault manager daemon (`fmfd`) can result in an `fmfd` core dump. The `fmfd` dumps core while it cleans up its resources, and this does not affect the FMA diagnosis.

Workaround: Add a few more strands into the primary domain. For example,

```
# ldm add-vcpu 3 primary
```

Attempting to WAN Boot a Logical Domain Using a Solaris 10 8/07 OS Installation DVD Causes Hang

Bug ID 6624950: WAN booting a logical domain using a miniroot created from a Solaris 10 8/07 OS installation DVD hangs during a boot of the miniroot.

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

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

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

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

Domain Services Uses One Lock for Read and Write, Which Can Cause a Logical Domain to Hang

Bug ID 6631043: This bug has not been seen on the Solaris OS. It has been seen on the virtual blade system controller (VBSC), which is running parallel code. It could cause the logical domain to hang.

Workaround: Stop and restart the affected logical domain.

Addition of Virtual Disk or Network Devices Under Delayed Reconfiguration Can Fail

Bug ID 6646690: If virtual devices are added to an active domain, and virtual devices are removed from that domain before that domain reboots, then the added devices do not function once the domain reboots.

Workaround: On an active domain, do not add and remove any virtual devices without an intervening reboot of the domain.

Recovery: Remove and then add the non-functional virtual devices, making sure that all the remove requests precede all the add requests, and then reboot the domain.

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

Bug ID 6656033: Simultaneous net installation of multiple guest domains fails on Sun SPARC Enterprise T5140 and Sun SPARC Enterprise T5240 systems that have a common console group.

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

System Can Panic During Reboot If the Virtual Switch Is Configured to Use an Aggregated Network Device

Bug ID 6678891: Occasionally, a service domain panics during reboot if the virtual switch is configured to use an aggregated network device for external connectivity.

Workaround: Configure the virtual switch to use a regular physical network device instead of an aggregated network device.

Recovery: Reconfigure the virtual switch to a physical network device using the `ldm set-vsw` command, and then restart the domain.

Sometimes, the `prtdiag(1M)` Command Does Not List All CPUs

Bug ID 6694939: In certain cases, the `prtdiag(1M)` command does not list all the CPUs.

Workaround: For an accurate count of CPUs, use the `psrinfo(1M)` command.

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

Bug ID 6687634: If the Sun Volume Manager (SVM) volume is built on top of a disk slice containing block 0 of the disk, then SVM prevents writing to block 0 of the volume to avoid overwriting the label of the disk.

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

Workaround: SVM volumes exported as a virtual disk should not be built on top of a disk slice containing block 0 of the disk.

A more generic guideline is that slices which start on the first block (block 0) of a physical disk should not be exported (either directly or indirectly) as a virtual disk. Refer to “Directly or Indirectly Exporting a Disk Slice” in the *Logical Domains (LDoms) 1.1 Administration Guide*.

Attempting a Network Interface Operation While Changing a Configuration Using CPR DR Can Cause the Logical Domains Manager to Terminate

Bug ID 6697096: For example, if you follow an `ldm rm-io` command by an `ldm set-vcpu` command, the Logical Domains Manager can, in certain circumstances, dump core and exit.

Workaround: For this specific example, reboot the domain after the `rm-io` subcommand and before the `set-vcpu` subcommand. In general, do not do a network interface operation while you are changing a configuration using CPU DR.

Recovery: The Service Management Facility (SMF) automatically restarts the Logical Domains Manager daemon (`ldmd`).

Deadlock Occurs Rarely With CPU DR Operations

Bug ID 6703958: Under rare circumstances, running CPU dynamic reconfiguration (DR) operations in parallel with network interface–related operations, such as `plumb` or `unplumb`, can result in a deadlock.

Workaround: Avoiding network interface–related operations can minimize this risk.

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

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

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

Cryptographic DR Changes Incompatible With Pre-LDoms Firmware

Bug ID 6713547: Cryptographic dynamic reconfiguration (DR) changes are incompatible with firmware that is prior to LDoms software releases. This problem prevents UltraSPARC T1-based systems running old firmware from using cryptographic hardware.

ZFS Pool Label Does Not Indicate That a Pool Is Closed Cleanly

Bug ID 6723511: The ZFS pool label does not indicate that a pool is closed cleanly. The following panic can occur if a disk image has been cloned from a guest domain with a different host ID.

```
misc/forthdebug (173689 bytes) loaded
WARNING: pool 'zfsroot' could not be loaded as it was last accessed by another
system (host:  hostid: 0x84a156d0).  See: http://www.sun.com/msg/ZFS-8000-EY

NOTICE:
spa_import_rootpool: error 9

Cannot mount root on /pci@400/pci@0/pci@1/scsi@0/disk@0,0:a fstype zfs

panic[cpu0]/thread=180e000: vfs_mountroot: cannot mount root

000000000180b940 genunix:vfs_mountroot+34c (800, 200, 0, 18c6400, 18f8000,
1921800)
  %10-3: 0000000001132400 0000000001132448 00000000018ce9b8 00000000012ae400
  %14-7: 00000000012ae400 0000000001923c00 0000000000000600 0000000000000200
000000000180ba00 genunix:main+120 (182f400, 191e400, 1870040, 1920c00, 180e000,
191a400)
  %10-3: 0000000001343800 000000000180bad0 0000000000004000 0000000001343800
  %14-7: 0000000000000000 000000000182f400 000000000182f628 0000000000000000
```

Workarounds: Use one of the following procedures:

- Boot into failsafe mode, which imports that pool with the correct host ID. Any subsequent reboots will work correctly.
- Boot the guest logical domain using a DVD, execute a `zfs import -f` command to change the ownership on the ZFS root pool (`rpool`) to the correct host ID, then reboot, and use the `rpool`.
- Reboot from the netinstall image, the miniroot, use the `zpool import -f` command to import the pool, and then immediately export the pool. Then reboot.

Sometimes CPU PM Fails to Retrieve PM Policy From the SP

Bug ID 6736962: CPU power management (PM) sometimes fails to retrieve the PM policy from the service processor (SP) when the Logical Domains Manager starts after the control domain boots.

If CPU power management could not retrieve the PM policy from the SP, then it allows the Logical Domains Manager to start as expected, but it does log the following error to the LDoms log and remains in performance mode:

```
Unable to get the initial PM Policy - timeout
```

Workaround:

1. Add the following line to the `/etc/system` file:

```
forceload: drv/ds_snmp
```

2. Reboot the control domain.

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

Bug ID 6742805: A domain shutdown or memory scrub can take over 15 minutes with 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 take quite a long time if a configuration is imbalanced; for example, a single CPU domain with 512GB of memory. This prolonged scrub time extends the amount of time it takes to shut down a domain.

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

Rarely, Control Domain Can Panic While Dynamically Removing a Virtual Network Interface From a Domain

Bug ID 6743338: Under rare circumstances, dynamically removing a virtual network interface from a domain can cause this domain to panic.

Workaround: Do not dynamically remove a virtual network interface just after it has been dynamically added to a domain, or just after the domain has booted.

Virtual Network Using Hybrid I/O Might Not Tag or Untag Packets When the Port VLAN ID Is Set

Bug ID 6746533: When the port VLAN ID (`pvid`) is set and hybrid I/O is enabled for a virtual network, the packets received and transmitted by the hybrid I/O resource to the outside network might not be tagged. Similarly, the received packets from the hybrid I/O resource might not be untagged before being sent up the stack.

Workaround: Do not enable hybrid I/O for a virtual network that has the `pvid` set.

Solaris 10/08 OS Installation Hangs With ZFS Boot on Certain Servers With 1-GB Memory

Bug ID 6747730: A Solaris 10/08 OS installation hangs with a ZFS boot on Sun SPARC Enterprise T5220 servers with 1-GB of memory.

Workaround: Perform the installation on a single disk, then establish the ZFS root mirror after rebooting.

Do Not Switch to Performance Mode in CPU Power Management Unless All Domains Are Up and Running the Solaris OS

Bug ID 6749619: Do not switch to performance mode in CPU power management unless all domains are up and running the Solaris OS. Otherwise, all the CPUs in guest domains might not be powered up and dynamically reconfigured.

Workaround: Before you switch to performance mode, check both domain and CPU power status by entering an `ldm list` command.

Recovery: If you are in a state where a guest domain has resources that did not power up while the system was in performance mode, toggle the policy to elastic mode and back to performance mode.

With Elara Copper Card, the Service Domain Hangs on Reboot

Bug ID 6753219: After adding virtual switches to the primary domain and rebooting, the primary domain hangs when installed with an Elara Copper card.

Workaround: Add this line to the `/etc/system` file on the service domain and reboot:

```
set vsw:vsw_setup_switching_boot_delay=300000000
```

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

Bug ID 6753683: Sometimes, executing the `uadmin 1 0` command from the command line of an LDom system does not leave the system at the OK prompt after the subsequent reset. This incorrect behavior is seen only when the LDom 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`.

Unrecoverable Hardware Error Seen on Guest With 2 Hybrid Virtual Networks When the Service Domain Stops

Bug ID 6756939: If a guest domain has 2 virtual networks enabled with hybrid I/O, and another guest domain configured as the service domain is stopped, an unrecoverable hardware error occurs.

Workarounds:

- Do not enable `set mode=hybrid` for more than one virtual network (`vnet`) in a guest domain.
- If you need to use more than one `vnet` in hybrid mode, then unplumb those `vnets` prior to rebooting the service domain that has the corresponding virtual switch (`vswitch`).

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

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

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

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

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

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

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

Newly Configured Virtual Network Fails to Establish a Connection With the Virtual Switch

Bug ID 6765355: Under rare conditions, when a new virtual network (`vnet`) is added to a logical domain, it fails to establish a connection with the virtual switch. This results in loss of network connectivity to and from the logical domain. If you encounter this error, you can see that the `/dev/vnetN` symbolic link for the virtual network instance is missing. If present, and not in error, the link points to a corresponding `/devices` entry as follows:

```
/dev/vnetN ->
../devices/virtual-devices@100/channel-devices@200/network@N:vnetN
```

Workarounds: Do one of the following:

- Perform a reconfiguration boot of the logical domain, whenever a `vnet` is added to the logical domain.
- If the logical domain is already booted, run the `devfsadm(1M)` command before plumbing the `vnet`.

Do Not Migrate a Guest Domain That Is at the `kmdb` Prompt

Bug ID 6766202: If a guest domain with only one CPU is at the kernel module debugger, `kmdb(1M)`, prompt, and if that domain is migrated to another system, the guest domain panics when it is resumed on the target system.

Workaround: Before migrating a guest domain, exit the `kmdb` shell, and resume the execution of the OS by typing `::cont`. Then migrate the guest domain. After the migration is completed, re-enter `kmdb` with the command `mdb -k`.

Cannot Export a ZFS Volume as a Single-Slice Virtual Disk From Service Domain Running Solaris 10 5/08 OS, or Earlier, to Guest Domain Running Solaris 10 10/08 OS

Bug ID 6769808: If a service domain running Solaris 10 5/08 OS, or earlier, is exporting a ZFS volume as a single-slice disk to a guest domain running Solaris 10 10/08 OS, then this guest domain is unable to use that virtual disk. Any access to the virtual disk fails with an I/O error.

Workaround: Upgrade the service domain to Solaris 10 10/08 OS.

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

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

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

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

Migration Does Not Fail If a `vdsdev` on the Target Has a Different Backend

Bug ID 6772120: If the virtual disk on the target machine does not point to the same disk backend that is used on the source machine, the migrated domain cannot access the virtual disk using that disk backend. 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 backends do not match.

Workaround: Ensure that when you are configuring the target domain to receive a migrated domain that the disk volume (`vdsdev`) matches the disk backend 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 backend:

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

After Node Reconfiguration, Logical Domains Manager Might Fail to Start on SPARC Enterprise T5440 Servers

Bug ID 6772485: The Logical Domains Manager daemon (`ldmd`) might fail to start after performing a node reconfiguration on a SPARC Enterprise T5440 server.

Workaround: If `CMP0` fails, move another Chip-Multithreaded Processor (CMP) module to CMP Slot 0, and reconfigure the system to operate in a degraded state. CMP Slot 0 always must be occupied by a working CMP module.

Constraint Database Is Not Synchronized to Saved Configuration

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

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

Cryptic Message to Indicate Incoming Migration Is Disabled on the Target

Bug ID 6773867: If the `incoming_migration_enabled` SMF property of the Logical Domains Manager daemon (`ldmd`) on a machine is set to `false` (by default it is `true`), and a user attempts to migrate a domain to the machine, the following cryptic message is printed by the Logical Domains Manager on the machine initiating the migration.

```
# ldm migrate-domain ldg1 target-machine
Target Password:
SSL ACCEPT FAILED ssl_err = 5 error:00000000:lib(0):func(0):reason(0)
Failed to connect to LDom manager on target-machine
Domain Migration of LDom ldg1 failed
```

Recovery: Set the `incoming_migration_enabled` property of the `svc:/ldoms/ldmd:default` SMF service back to `true` using the `svccfg(1M)` command.

Two Faulty CPUs That Include CPU 0 Can Cause Power Management Utilization Issue After Rebooting

Bug IDs 6773930 and 6779134: You might experience this problem if the system's CPU power management (PM) policy is set to `elastic`, there is more than one faulty CPU, and one of the faulty CPUs is `cpu0`. You can see this by using the `psrinfo` or `fmadm faulty` commands.

Workaround: Switch the power management policy on the service processor (SP) to performance, and then back to elastic.

```
-> set SP/powermgmt policy=performance
-> set SP/powermgmt policy=elastic
```

Domain Can Lose Virtual CPUs During a Migration If Another Domain Is Rebooting

Bug ID 6775847: There is a period of time where a domain being migrated onto a system can end up with just one virtual CPU or hung if another domain on the target system is rebooted during the migration. The `start-domain` and `stop-domain` operations of the `ldm(1M)` command are prevented currently, but issuing a `reboot` and `init` command in the Solaris OS instance running on a guest domain cannot be prevented.

Workaround: Avoid rebooting domains while a migration is in progress onto a machine.

Recovery: Stop and restart the migrated domain on the target system if you detect the symptoms of this issue.

Panic Occurs When More Than Two Hybrid I/O-Capable Virtual Networks Are Activated on a Guest Domain

Bug ID 6777756: A panic occurs when more than two Hybrid I/O-capable virtual networks are activated in a guest domain.

Recovery: Remove all entries in `/etc/path_to_inst` of the guest domain that are similar to the following, and reboot:

```
"/niu@80/network@bad1" 2 "nxge"
"/niu@80/network@bad1" 3 "nxge"
"/niu@80/network@bad1" 4 "nxge"
"/niu@80/network@bad1" 5 "nxge"
```

Only entries, such as those following, that have function number 0 and 1 are known to not create this issue.

```
"/niu@80/network@bad1" 0 "nxge"
"/niu@80/network@bad5" 1 "nxge"
```


Migration Does Not Clean Up on Target If Virtual Network MAC Address Clashes With Existing Domain

Bug ID 6779482: If a domain being migrated has a virtual network (`vnet`) with a MAC address that matches a MAC address on the target, the migration fails appropriately, but leaves a residual inactive domain of the same name and configuration on the target.

Workaround: On the target, use the `ldm destroy` command to remove the inactive domain manually, fix the MAC address so that it is unique, and try the migration again.

Explicit Console Group and Port Bindings Are Not Migrated

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

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

Migration Dry Run Does Not Detect Inadequate Memory

Bug ID 6783450: The domain migration dry run option (`-n`) does not ensure there is enough memory free on the target system to bind the domain specified. If all other criteria are met, the command will return without an error but will correctly return an error when the migration is actually attempted.

Workaround: Run the `ldm list-devices mem` command on the target machine to verify that there is enough memory available for the domain to be migrated.

In Certain Situations, a Domain Migration Can Fail After the Control Domain Is Rebooted

Bug ID 6784943: After the control domain has been rebooted, the first migration attempt might fail with the following error message:

```
Failed to send 'migrate' command to ldm(1m)
```

This occurs if the Logical Domains Manager is started at boot time after networking is initialized. This includes scenarios where the control domain is booted without a plumbed network device and where a network device is set up after the boot has completed.

Workaround: Restart the Logical Domains Manager once networking is active in the control domain.

```
# svcadm restart ldmd
```

Recovery: Once this problem occurs, the Logical Domains Manager is restarted automatically, clearing the error condition. Future attempts to perform migration should be successful.

Pseudonyms for PCI Busses on Sun SPARC Enterprise T5440 Systems Are Not Correct

Bug ID 6784945: On a Sun SPARC Enterprise T5440 system, the pseudonyms (shortcut names) for the PCI busses are not correct.

Workaround: To configure PCI busses on a Sun SPARC Enterprise T5440 system, you must use the `pci@xxx` form of the bus name, as listed under the `DEVICE` column of any of the following list commands:

- `ldm ls -l ldom`
- `ldm ls -o physio ldom`
- `ldm ls-devices`

Cancelling Domain Migration With Virtual Networks Using Multiple Virtual Switches Might Cause Domain Reboot

Bug ID 6787057: On a guest domain with two or more virtual network devices (`vnets`) using multiple virtual switches (`vsws`), if an in-progress migration is cancelled, the domain being migrated might reboot instead of resuming operation on the source machine with the OS running. This issue does not occur if all the `vnets` are connected to a single `vsw`.

Workaround: If you are migrating a domain with two or more virtual networks using multiple virtual switches, do not cancel the domain migration (either using `^C` or the `ldm cancel-operation` command) once the operation starts. If a domain is inadvertently migrated, it can be migrated back to the source machine once the original migration is completed.

Possibility That New Domain on Source System Could Be Assigned Same MAC Address as Domain That Was Successfully Migrated

Bug ID 6788178: After a domain has successfully migrated from a source system to a target system, it is possible that a newly created domain on the source system could be allocated the same MAC address as the domain that was successfully migrated. If the source and target systems are on the same subnet, this can result in the new domain being unable to communicate on the network. In this case, the Solaris OS might generate errors stating that another machine on the network has the same IP address.

Workaround: If this issue occurs, operation can be restored by changing the MAC address of virtual network interfaces having problems.

▼ Change the MAC Address of an Affected Virtual Network Interface

1. Stop the affected domain (for example, `ldg1`).

```
# ldm stop ldg1
```

2. Unbind the affected domain.

```
# ldm unbind ldg1
```

3. Change the MAC address of the affected virtual network interface (for example, `vnetX`).

```
# ldm set-vnet mac-addr=xx:xx:xx:xx:xx:xx vnetX ldg1
```

4. Bind the domain.

```
# ldm bind ldg1
```

5. Start the domain.

```
# ldm start ldg1
```

Documentation Errata

This section contains documentation errors that have been found too late to resolve for the LDomS 1.1 release.

VIO DR Operations Ignore the Force (-f) Option

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

The `ldm rm-reconf` Command No Longer Works

Bug ID 6774570: The `ldm` man page and the *Logical Domains Manager (LDoms) 1.1 Man Page Guide* erroneously state that you can still use the `ldm rm-reconf` command as an alias for the new `ldm cancel-operation reconf` command.

Workaround: Use the `ldm cancel-operation reconf` command to cancel a delayed reconfiguration operation.

The `vntsd(1M)` Man Page Description for `listen_addr` Is Incomplete

The revised portion, which is not in the Solaris 10 10/08 Reference Manual Collection, now reads:

```
vntsd/listen_addr
```

Set the IP address to which `vntsd` listens, using the following syntax:

```
vntsd/listen_addr:"xxx.xxx.xxx.xxx"
```

...where `xxx.xxx.xxx.xxx` is a valid IP address. The default value of this property is to listen on IP address `127.0.0.1`. Users can connect to a guest console over a network if the value is set to the IP address of the control domain.

Note -

Enabling network access to a console has security implications. Any user can connect to a console and for this reason it is disabled by default.

Error in Descriptions of `default-vlan-id`, `pvid`, and `vid` in the `ldm(1M)` Man Page

In the `ldm(1M)` man page in the section on using the `add-vsw` subcommand, the definitions of `default-vlan-id`, `pvid`, and `vid` should say:

- `default-vlan-id=vlan-id` specifies the default virtual local area network (VLAN) to which a virtual switch and its associated virtual network devices belong to implicitly, in untagged mode. It serves as the default port VLAN id (`pvid`) of the virtual switch and virtual network devices. Without this option, the default value of this property is 1. Normally, you would not need to use this option. It is provided only as a way to change the default value of 1.
- `pvid=port-vlan-id` specifies the VLAN to which the virtual switch needs to be a member, in untagged mode. This applies to the `set-vsw` subcommand also.
- `vid=vlan-id` specifies one or more VLANs to which a virtual switch needs to be a member, in tagged mode. This applies to the `set-vsw` subcommand also.

In the `ldm(1M)` man page in the sections on using the `add-vnet` and `set-vnet` subcommands, the definitions of `pvid` and `vid` should say:

- `pvid=port-vlan-id` specifies the VLAN to which the virtual network device needs to be a member, in untagged mode.
- `vid=vlan-id` specifies one or more VLANs to which a virtual network device needs to be a member, in tagged mode.

The `status` Output Option Was Omitted for the `list` Subcommand in the `ldm(1M)` Man Page

The `list` subcommand has new output (`-o format`) options to limit the output format. The `status` output option was omitted from output options available in the `ldm(1M)` man page. This option is used to check the status of a migrating domain.

Resolved Issues

This section contains bugs that have been fixed since the LDomS 1.0.3 software release.

LDoms 1.1 RFEs and Bugs Fixed in Solaris 10 10/08 OS

The following LDomS requests for enhancements (RFEs) and bugs were fixed for the Solaris 10 10/08 OS release:

TABLE 6 LDomS 1.1 RFEs and Bugs Fixed in Solaris 10 10/08 OS

Bug ID	Synopsis
6405398	LDoms NIU Hybrid I/O support
6411419	LDC ACK packets can be blocked by unread data packets
6419257	<code>ldc_chkq</code> incorrectly returns that there is data available to read when there is not
6434157	<code>libefi</code> sometimes gets the wrong <code>gpt</code> partition location
6503157	Guest domain I/O statistics need to be made available to <code>iostat(1M)</code>
6528974	Usage of <code>DKIOCGETEFI</code> has changed and impacts <code>vdisk</code> s using ZFS volumes
6533308	File management daemon (<code>fmd</code>) does not properly recover from an LDC channel reset
6552999	If you press <code>Control-c</code> from the <code>prtdiag</code> command, you receive blank environmental data fields when you run again
6558966	Virtual disks created from files do not support EFI labels
6560890	DS error handling and error/debug messaging needs tidying up
6561424	Heavy network traffic in LDomS guest domains leads to Solaris Cluster heartbeat failures
6569471	<code>VIO_IS_VALID_DESC_STATE</code> and <code>VIO_SET_DESC_STATE</code> utterly broken (but unused)
6581309	Inconsistent console behavior when not using the virtual console
6581655	Stream mode LDC channels are not properly reset
6610700	<code>vswitch</code> should be selective when printing notices and warnings to console
6621749	Cannot add LDomS virtual disks with a VTOC label to an OBAN disk set
6622004	VDC should support multipathing using multiple <code>vdisk</code> servers
6635697	QCN driver should support a polled I/O mode

TABLE 6 LDoms 1.1 RFEs and Bugs Fixed in Solaris 10 10/08 OS (*Continued*)

Bug ID	Synopsis
6635768	vnet and vsw should support VLANs
6637560	Disks are not correctly exported with vxdump
6637596	Invalid assertion in ip_soft_ring_assignment()
6638558	Memory leak on ddi_prop_decode_alloc+0xc when running SunVTS
6643903	Possible memory leak in vsw driver
6651197	Add support for LDoms Virtual I/O Dynamic Reconfiguration (VIO DR)
6653726	Guest domain panics on page_get_replacement_page in sparse-memory, memory-exhaustion test case
6666476	LDC mode names do not conform to FWARC 2006/571
6671853	cnex interrupt distribution is imbalanced
6673364	VDC should support DKIOCPARTITION ioctl
6675762	VDC should check for device being read-only on open
6675887	No vsw-to-vsw connectivity on vsw services bound to an aggregated device
6680459	VDC should define Nblocks/Size as dynamic properties
6682775	VDS code duplicates is_pseudo_device function
6685162	Solaris OS cannot be installed on a single-slice disk
6687871	SPARC Enterprise T2000 system hangs on boot due to vsw issues with nxge breakage
6688984	VDC race condition seems to be using freed buffer from vdc_process_data_msg
6689871	Lock in vsw_tx_msg inhibits guest network throughput
6690439	Deadlock during attach when a domain exports a vdisk to itself
6690911	vdisk read/write with a size which is not a multiple of 8 bytes hangs
6694540	Disks are not correctly exported with EMC power path
6695641	Assert failure in vds:vd_recv_msg() when resuming a suspended domain
6704840	Unable to use a 300G zvol as backend to a vdisk in an ldom guest
6706799	vsw needs to try cleaning up shares at reboot or panic time
6713652	The ldm set-vcpu 1 primary command causes hang
6715584	The vdc_strategy should use polled I/O during panic
6717698	Guest domain panics at n2cp:block_final_start during PCK11.auto test run
6719086	Dynamic update of net-dev in vsw makes guests inaccessible for a short time

RFEs and Bugs Fixed for LDom 1.1 Software

The following LDom 1.1 RFEs and bugs were fixed for the LDom 1.1 software release:

TABLE 7 RFEs and Bugs Fixed for LDom 1.1 Software

Bug ID	Synopsis
6510365	<code>ldmd</code> needs to support key store domain service
6519049	<code>add-vsw</code> returns misleading message if <code>ldom</code> or <code>vswitch_name</code> arguments are missing
6542175	LDom Manager lacks bounds checking for RA segment IDs when calculating RA values
6562222	<code>ldm list</code> CLI should provide a means to request a subset of the full listing
6567372	<code>ldmd</code> should provide the highest numbered <code>vcpu</code> for DR <code>rm-vcpu</code> command
6573220	<code>ldm add-vcc</code> , wrong error message displayed when command fails
6586046	Allow <code>ldm list -l ldom</code> to display MAC assigned to the domain
6590124	Oddly worded error message when trying to delete a bound domain
6591905	<code>ldm list -l</code> reports duplicate entry of NIU
6596652	Upon restart, the Logical Domains Manager reverts automatic console port selection on bound guests into specific port constraint
6597377	<code>ldm set-vcons</code> command should require port and group arguments
6611266	<code>ldmd -g -p</code> with no arguments causes message "Segmentation Fault - core dumped"
6616398	<code>ldmd -p invalid port number</code> does not return an error as expected
6646277	<code>ldm -V</code> option is not listed in the usage messages displayed
6655083	<code>ldm panic-domain</code> returns incorrect error message
6663457	<code>ldm ls-services</code> CLI should list domain to which the service belongs
6665919	Add support for LDom Virtual I/O Dynamic Reconfiguration (VIO DR)
6670605	Want an interface to set LDom <code>hostid</code>
6674948	SUNWldm postinstall script does not handle alternate root paths well
6682951	Restore original error message to show guest domain bind failure due to non-supported CONS operation
6683498	LDom Manager support for NIU Hybrid I/O
6685149	Lost <code>devdisk vdisks</code> after <code>ldm rm-vds primary-vds0 primary</code> command and <code>ldm cancel-reconf primary</code> command
6686963	Add LDom Manager support for VLAN tagging
6688407	Upon enabling <code>ldmd</code> service, warning "No address-mask property in dma-latency-group node" seen in log

TABLE 7 RFEs and Bugs Fixed for LDomS 1.1 Software (*Continued*)

Bug ID	Synopsis
6689040	ldmd returns "0" on two negative test cases: "ldm rm-vcpu" returns 0 when CPU should not be removed
6692185	rm-vdsdev succeeds on another domain while there are delayed reconfig operations pending
6692261	Need XML support for new VLAN parameters
6693542	When vdsdev and vdisk are in the same domain, rm-vdsdev does not persist across reboots
6695424	set-vconsole CLI not assigning the group name and port number for already defined service
6696250	Need better handling of error from pri_get (PRI_WAITGET) calls
6696286	CLI list-constraints needs to show the VSW mode
6697873	Existing vswitch MAC address is lost if a delayed reconfiguration is cancelled after a set-vsw
6697940	set-vcc does not roll back to prior port-range value on delayed reconfiguration cancellation
6699332	LDoms Manager should reserve real address range for mapping remote memory
6700002	set-vsw should not allow change of MAC address if the domain is bound or active
6700129	Received 00:00:7f:ff:ff:ff as MAC address for all its VIO devices after rebinding of an I/O domain
6708814	Identify I/O domains in mdstore domain service when storing to vbosc
6709020	LDoms Manager getting incorrect vcpu utilization from hypervisor
6711897	hostid is not displayed for guest domains unless explicitly assigned during creation
6711904	Manually assigned MAC addresses should be allowed to be assigned in multiples
6718108	LDoms uses wrong hostid prefix
6724210	ldmd return incorrect error messages from DR mau(s)
6726072	CPU DR should not partially add CPUs to the guest domain during guest domain boot time
6726177	Need a meaningful error for failing CPU DR operations to a guest domain in the process of booting
6727074	XMLv3: add-crypto/set-crypto/remove-crypto action tags fail with: "not a known action tag"
6727293	ldm list-constraints -x of a non-existent domain is handled incorrectly
6729544	VIO nodes in guest MD needs additional properties to recover previous state without constraint database
6730974	Automatic allocating of the domain MAC address does not save the MAC address in the MAC address table
6732861	ldm ls -l shows OpenBoot PROM security password in clear text for primary and guest domains
6734471	Restarting ldmd means hostid in inactive domain gets set to 0
6737032	Logical Domains Manager memory leak plus a reuse of freed memory

TABLE 7 RFEs and Bugs Fixed for LDomS 1.1 Software (*Continued*)

Bug ID	Synopsis
6740525	XMPP server should fail if a client sends XML that is not expected
6741733	LDoms Manager should display better error messages for DR CPU failures
6743537	Non-debug package should not be built with debugging macros, etc.
6744046	<code>ldm list-config</code> output can be confusing when issued after a delayed reconfiguration reset
6756001	Cryptographic (mau) operations are not treated as delayed reconfiguration by default on guest domains
6757682	MAC multicast code allows negative numbers for the <code>hops</code> (TTL) variable
6769790	If <code>hostent</code> is null, <code>ldmd</code> crashes in <code>inet_ntop()</code> call
6783465	<code>ldmd</code> SMF manifest start/stop values need to be bumped up for power management (PM)