



# Configuring logical domains and SFHA in OVM for SPARC environment with multiple IO Domains

Sajith C R, Senior Software Engineer  
Information Availability Group

# Content

- Introduction ..... 3**
- Advantage of multiple IO domains ..... 3
- Configuring logical domains and SFHA in OVM for SPARC environment with multiple io domains..... 4**

## Introduction

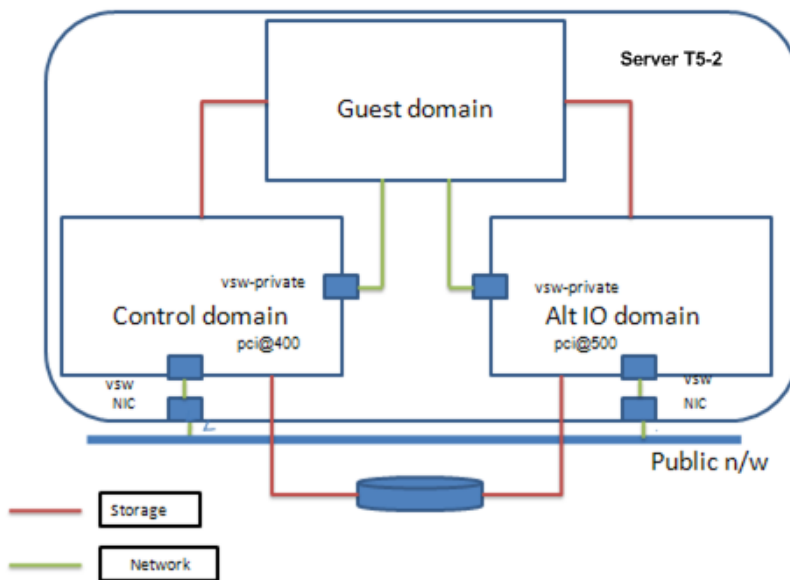
This white paper describes how to configure Oracle VM for SPARC logical domains with multiple I/O domains with Symantec Storage Foundation and High Availability (SFHA).

The logical domains are managed by Symantec Cluster Server (VCS) installed in the I/O domains and the logical domains are part of a multi-node SFHA cluster formed between logical domains.

This whitepaper starts with converting an existing physical system to support logical domains, describes the steps to export storage and network services to logical domains, and provides steps to install SFHA or VCS in logical domains and I/O domains.

### Advantage of multiple IO domains

If the physical system has multiple host bus adapters (HBA) and NICs from different PCI buses, you can configure multiple I/O domains. Multiple I/O domains ensure that logical domains get storage and network services from more than one I/O domain. This helps logical domain to have continuous access to storage and network connectivity in the event of planned or unplanned outage of one of the I/O domain.



## Configuring logical domains and SFHA in OVM for SPARC environment with multiple io domains

Perform the following steps:

1. Install Oracle VM Server for SPARC software and create default services for primary domain.

On Oracle Solaris 11, OVM for SPARC package (`system/ldoms/ldomsmanager`) comes with OS installation. On Oracle Solaris 10, you have to download the required version of `SUNWldomr` and `SUNWldomu` packages and install them.

Ensure that you have the latest version of OVM for SPARC and corresponding firmware supported by VCS installed on your system.

Refer to the Oracle VM Server for SPARC administration guide for install and upgrade steps.

After you install the OVM Server for SPARC, create the default services for virtual console concentrator service, virtual disk server, and virtual switch service.

```
primary# ldm add-vcc port-range=5000-5100 primary-vcc0 primary
```

```
primary# ldm add-vds primary-vds0 primary
```

```
primary# ldm add-vsw net-dev=nxge0 primary-vsw0 primary
```

(On Solaris 11, use `net0` as network adapter driver)

2. Create an alternate I/O domain.

Refer to the Oracle VM Server for SPARC administration guide for steps to set up alternate I/O domain.

You may create more than one IO domain in addition to the primary domain depending upon the specification of your server.

For example, you may create alternate io domain with `pci_2` bus.

```
primary# ldm add-domain alternate
```

```
primary# ldm list-io
```

```
primary# ldm remove-io pci_2 primary
```

```
primary# ldm add-config multiple-io1
```

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

```
primary# ldm add-io pci_2 alternate
```

```
primary# ldm set-mem 4G alternate
```

```
primary# ldm set-vcpu 16 alternate
```

```
primary# ldm bind alternate
```

```
primary# ldm start alternate
```

```
primary# ldm add-config multiple-io2
```

Set up services for virtual console concentrator service, virtual disk server, and virtual switch service for alternate IO domain.

```
primary# ldm add-vcc port-range=5000-5100 alternate-vcc0 alternate
```

```
primary# ldm add-vds alternate-vds0 alternate
```

```
primary# ldm add-vsw net-dev=nxge0 alternate-vsw0 alternate
```

3. Create a guest domain.

Create a guest domain with adequate vcpu and memory.

```
primary# ldm add-domain ldm1
```

```
primary# ldm set-mem 4G ldm1
```

```
primary# ldm set-vcpu 16 ldm1
```

4. Ensure that storage requirements are met.

While planning for creation of logical domain, you also need to identify the storage required for the application data and OS image. In addition, if you are planning to use disk based I/O fencing, you need to assign SCSI3 compliant disks of minimum 3 disks (odd numbered) with minimum size 256MB.

After you identify the storage requirements, assign the virtual disk (vdisk) from all the I/O domains to the logical domain.

You can choose one of the following methods while assigning virtual disks to logical domain:

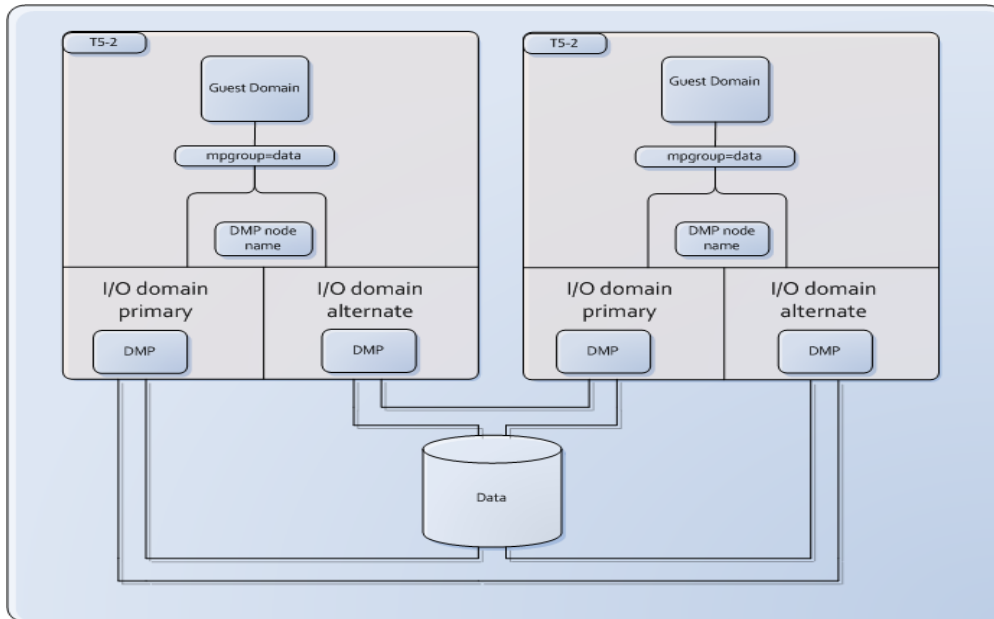
I) Using virtual disk multipathing (MPGROUP)

**Note:** If the `mpgroup` attribute is set while exporting virtual disk to logical domain, you cannot use disk based (SCSI3-PGR) I/O fencing inside the logical domain.

- a) Virtual disk multipathing enables you to configure a virtual disk on a guest domain to access its back-end storage by more than one path. The paths lead through different service domains that provide access to the same back-end storage, such as a disk LUN. This feature enables a virtual disk in a guest domain to remain accessible even if one of the service domains goes down. To enable virtual disk multipathing, you must export a virtual disk back end from each service domain and add the virtual disk to the same multipathing group (`mpgroup`).

With this method, define virtual disk server device (`vdsdev`) for the storage path from I/O domain with the `mpgroup` attribute set. Similarly, define the virtual disk server device for the same storage device from other I/O domain with the same `mpgroup` attribute value. Then, export the virtual disk from one of the I/O domain.

- b) If you have multiple HBAs available from each I/O domain and Symantec Dynamic Multipathing (DMP) is installed on each I/O domain, you may export the storage path using DMP meta-node so that you can avoid exporting multiple virtual disk to logical domain.



For example, to export boot disk to logical domain run the following commands:

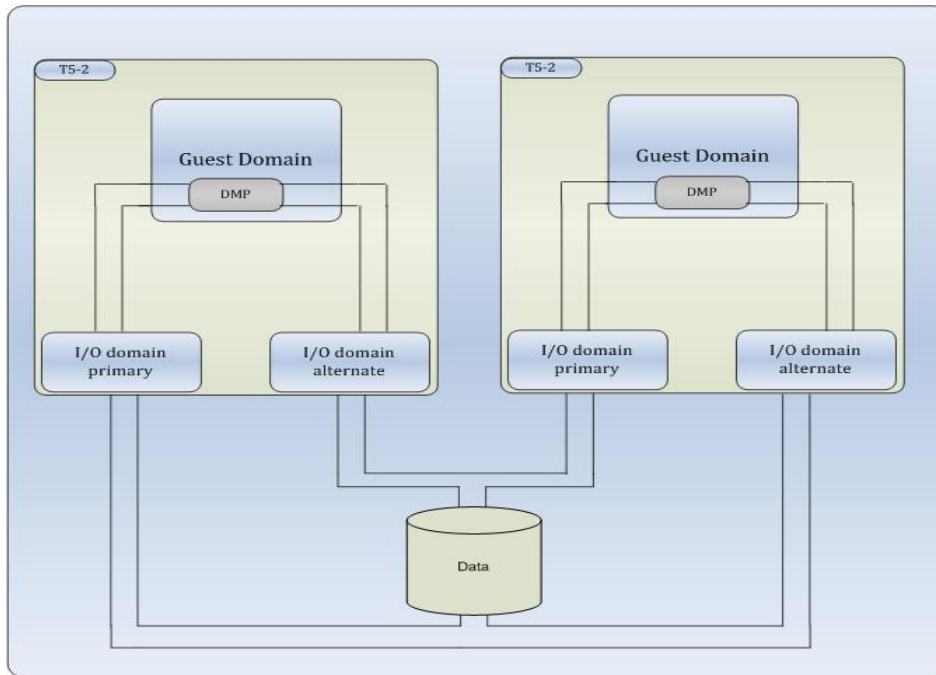
```
primary# ldm add-vdsdev mpgroup=ldm1_boot /dev/dsk/c0t5006048C53694120d1826s2
ldm1_boot@primary-vds0
primary# ldm add-vdsdev mpgroup=ldm1_boot /dev/dsk/c1t5006048C53694120d1826s2
ldm1_boot@alternate-vds0
primary# ldm add-vdisk timeout=10 ldm1_boot ldm1_boot@primary-vds0 ldm1
```

**Note:** Do not set the value of the Options attribute to exclusive; excl. If Options attribute is set to exclusive, logical domain cannot use the multipathing functionality.

Repeat the steps for each disk you want to export to the logical domain.

## II) Using DMP at the guest domain for multipathing.

- a) With this method, you can export either all the storage path from each I/O domains to the logical domain individually and use DMP inside the logical domain for multipathing.
- b) If you have DMP version higher than 5.1.SP1 installed in I/O domains, you can export DMP meta-node to logical domain so that you can avoid exporting multiple virtual disk paths to logical domain. If you are planning to use SCSI3-PGR based I/O fencing inside the logical domains, you need to have DMP version higher than 6.1.0 HF1 in the I/O domains.



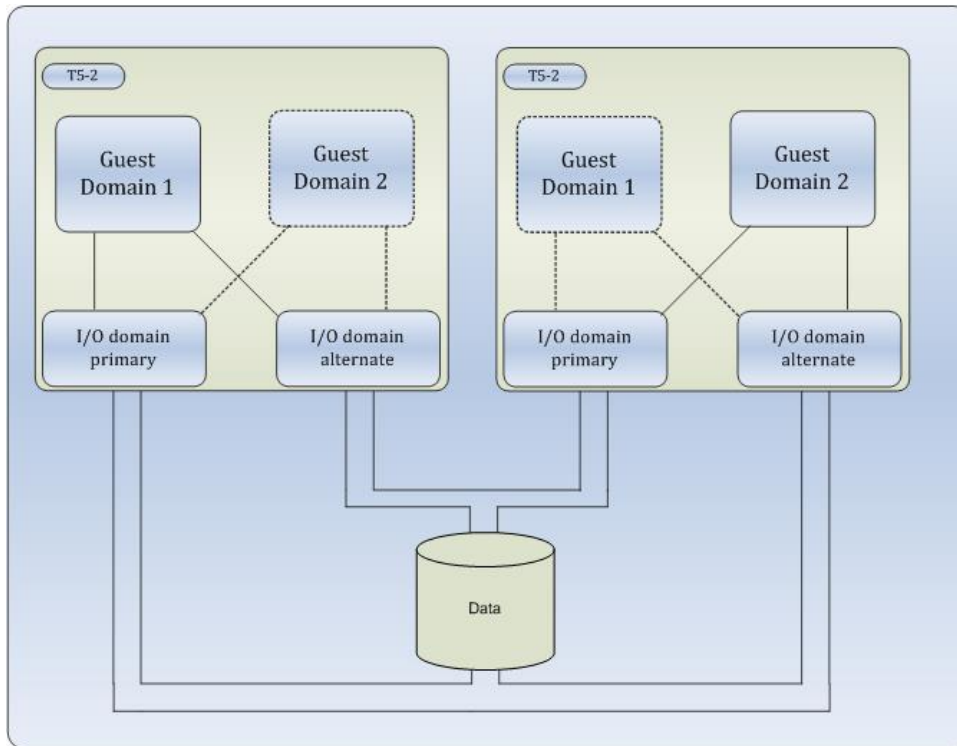
For example, to export boot disk to logical domain run the following commands:

```
primary# ldm add-vdsdev /dev/vx/dmp/emc_clariion0_1 ldm1_boot@primary-vds0
primary# ldm add-vdsdev /dev/vx/dmp/emc_clariion0_1 ldm1_boot@alternate-vds0
primary# ldm add-vdisk timeout=10 ldm1_boot ldm1_boot@primary-vds0 ldm1
```

Repeat the steps for each disk you want to export to the logical domain.

**Note:**

1. Symantec does not support I/O fencing to be enabled on both the control and logical domain clusters.
2. If you have plan to use disk based I/O fencing inside logical domain and use the OVM live migration feature with the logical domains, make sure that the storage path exported to the logical domain is unique on the server. For example, two logical domains should not use same storage path when running on the same physical server. So, if you have more than one HBA in an I/O domain, you can opt to provide storage device path from one HBA for one of the logical domain and storage device path from another HBA for the other logical domain. Also DMP version on the logical domain must be at least 6.0.5.
3. DMP support for ZFS root pool is available form version 6.1.0 onwards.



5. Ensure that network requirements are met.

Assign virtual network interfaces from multiple I/O domains to the logical domain.

```
primary# ldm add-vnet linkprop=phys-state id=0 vnet1 primary-vsw0 ldm1
primary# ldm add-vnet linkprop=phys-state id=1 vnet2 alternate-vsw0 ldm1
```

Repeat the steps for each virtual nic that you want to export to the logical domain.

Export enough virtual NICs to logical domains considering application requirement and VCS interconnect requirements. While exporting the virtual NICs, ensure that for virtual NICs meant for public network, the back-end physical NIC on I/O domain is connected to public network. Similarly, ensure that virtual NICs meant for cluster interconnects are physically connected across I/O domains. Inside the guest domain, after you install the OS, you may use IPMP or MultiNIC agents from VCS to manage the virtual NICs exported from I/O domains.

***Perform the steps from 1 to 5 on all physical systems where you want to create logical domains.***

6. Install VCS on all I/O domains.

Download and install the VCS product on all I/O domains. If you want to use DMP, install DMP also on all I/O domains. Refer to the product install guide for installation details.

Configure resources for LDom, AlternateIO, Storage (e.g: Zpool, Disk) and Network (e.g: NIC). Refer to the Storage Foundation Virtualization guide for Solaris for more details.

A sample cluster configuration with multiple I/O domains is provided in the appendix.



7. Install SFHA on logical domains.

Download and install the SFHA product on the logical domains.

Refer to the product install guide for steps to install SFHA.

**Reference:**

You can refer to Symantec Storage Foundation and High Availability Solutions Virtualization Guide on Solaris for details on supported configuration

## APPENDIX

See the Sample main.cf for a cluster between two physical systems with primary and alternate I/O domain.

Logical domain is configured with storage and network services from primary domain and alternate I/O domain.

```
include "OracleASMTypes.cf"
include "types.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.cf"

cluster vcs-demo-cluster (
    UserNames = { admin = cNOgNInKOjOOmWOiNL }
    Administrators = { admin }
)
system primary1 (
)
system alternate1 (
)
system primary2 (
)
system alternate2 (
)
group AltIO_group1 (
    SystemList = { primary1 = 0, primary2 = 1 }
    Parallel = 1
)
AlternateIO AlternateIO_res1 (
    StorageSG @primary1 = { Storage_group_ldm1_dom1 = 0 }
    StorageSG @primary2 = { Storage_group_ldm1_dom2 = 0 }
    NetworkSG @primary1 = { Network_group_ldm1_dom1 = 0 }
    NetworkSG @primary2 = { Network_group_ldm1_dom2 = 0 }
)
// resource dependency tree
//
// group AltIO_group1
// {
// AlternateIO AlternateIO_res1
// }
```

```

group AltIO_group2 (
    SystemList = { primary1 = 0, primary2 = 1 }
    Parallel = 1
)
AlternateIO AlternateIO_res2 (
    StorageSG @primary1 = { Storage_group_ldm2_dom1 = 0 }
    StorageSG @primary2 = { Storage_group_ldm2_dom2 = 0 }
    NetworkSG @primary1 = { Network_group_ldm2_dom1 = 0 }
    NetworkSG @primary2 = { Network_group_ldm2_dom2 = 0 }
)
// resource dependency tree
//
// group AltIO_group2
// {
// AlternateIO AlternateIO_res2
// }

```

```

group LDom_group1 (
    SystemList = { primary1 = 0, primary2 = 1 }
)
LDom LDom_ldm1 (
    LDomName = ldm1
    CfgFile = "/LDomCfg/ldm1.xml"
    RemoveLDomConfigForMigration = 1
    Password = iumUnuVuxOjoKol
)
requires group AltIO_group1 online local hard
// resource dependency tree
//
// group LDom_group1
// {
// LDom LDom_ldm1
// }

```

```

group LDom_group2 (
    SystemList = { primary1 = 0, primary2 = 1 }
)
LDom LDom_ldm2 (
    LDomName = ldm2
    CfgFile = "/LDomCfg/ldm2.xml"
    RemoveLDomConfigForMigration = 1
)

```

```

        Password = iumUnuVuxOjoKol
    )
requires group AltIO_group2 online local hard
// resource dependency tree
//
//   group LDom_group2
//   {
//   LDom LDom_ldm2
//   }
group Network_group_ldm1_dom1 (
    SystemList = { alternate1 = 1, primary1 = 0 }
    Parallel = 1
)
NIC nicres1 (
    Device @alternate1 = nxge0
    Device @primary1 = nxge3
)
Phantom phres2 (
)
// resource dependency tree
//
//   group Network_group_ldm1_dom1
//   {
//   NIC nicres1
//   Phantom phres2
//   }
group Network_group_ldm1_dom2 (
    SystemList = { primary2 = 0, alternate2 = 1 }
    Parallel = 1
)
NIC nicres2 (
    Device @primary2 = nxge3
    Device @alternate2 = nxge0
)
Phantom phres4 (
)
// resource dependency tree
//
//   group Network_group_ldm1_dom2
//   {

```

```

// NIC nicres2
// Phantom phres4
// }
group Network_group_ldm2_dom1 (
  SystemList = { alternate1 = 1, primary1 = 0 }
  Parallel = 1
)
NIC nicres3 (
  Device @alternate1 = nxge0
  Device @primary1 = nxge3
)
Phantom phres5 (
)
// resource dependency tree
//
// group Network_group_ldm2_dom1
// {
//   NIC nicres3
//   Phantom phres5
// }
group Network_group_ldm2_dom2 (
  SystemList = { primary2 = 0, alternate2 = 1 }
  Parallel = 1
)
NIC nicres4 (
  Device @primary2 = nxge3
  Device @alternate2 = nxge0
)
Phantom phres6 (
)
// resource dependency tree
//
// group Network_group_ldm2_dom2
// {
//   NIC nicres4
//   Phantom phres6
// }
group Storage_group_ldm1_dom1 (
  SystemList = { primary1 = 0, alternate1 = 1 }
  Parallel = 1

```

```
)  
Disk diskres1 (  
    Partition @primary1 = "/dev/rdisk/c2t50060E8006D43C01d2s2"  
    Partition @alternate1 = "/dev/rdisk/c0t50060E8006D43C01d2s2"  
)  
Disk diskres2 (  
    Partition @primary1 = "/dev/rdisk/c2t50060E8005650921d0s2"  
    Partition @alternate1 = "/dev/rdisk/c0t50060E8005650921d0s2"  
)  
Disk diskres3 (  
    Partition @primary1 = "/dev/rdisk/c2t50060E8005650921d1s2"  
    Partition @alternate1 = "/dev/rdisk/c0t50060E8005650921d1s2"  
)  
Phantom phres1 (  
)  
// resource dependency tree  
//  
// group Storage_group_ldm1_dom1  
// {  
// Disk diskres1  
// Disk diskres2  
// Disk diskres3  
// Phantom phres1  
// }  
group Storage_group_ldm1_dom2 (  
    SystemList = { primary2 = 0, alternate2 = 1 }  
    Parallel = 1  
)  
Disk diskres4 (  
    Partition @primary2 = "/dev/rdisk/c2t50060E8006D43C01d2s2"  
    Partition @alternate2 = "/dev/rdisk/c0t50060E8006D43C01d2s2"  
)  
Disk diskres5 (  
    Partition @primary2 = "/dev/rdisk/c2t50060E8005650921d0s2"  
    Partition @alternate2 = "/dev/rdisk/c0t50060E8005650921d0s2"  
)  
Disk diskres6 (  
    Partition @primary2 = "/dev/rdisk/c2t50060E8005650921d1s2"  
    Partition @alternate2 = "/dev/rdisk/c1t50060E8005650921d1s2"  
)
```

```

Phantom phres3 (
    )
// resource dependency tree
//
//  group Storage_group_ldm1_dom2
//  {
//  Disk diskres4
//  Disk diskres5
//  Disk diskres6
//  Phantom phres3
//  }
group Storage_group_ldm2_dom1 (
    SystemList = { primary1 = 0, alternate1 = 1 }
    Parallel = 1
)
Disk diskres7 (
    Partition @primary1 = "/dev/rdisk/c2t50060E8006D43C01d3s2"
    Partition @alternate1 = "/dev/rdisk/c0t50060E8006D43C01d3s2"
)
Disk diskres8 (
    Partition @primary1 = "/dev/rdisk/c2t50060E8005650921d0s2"
    Partition @alternate1 = "/dev/rdisk/c0t50060E8005650921d0s2"
)
Disk diskres9 (
    Partition @primary1 = "/dev/rdisk/c2t50060E8005650921d1s2"
    Partition @alternate1 = "/dev/rdisk/c0t50060E8005650921d1s2"
)
Phantom phres7 (
    )
// resource dependency tree
//
//  group Storage_group_ldm2_dom1
//  {
//  Disk diskres7
//  Disk diskres8
//  Disk diskres9
//  Phantom phres7
//  }
group Storage_group_ldm2_dom2 (
    SystemList = { primary2 = 0, alternate2 = 1 }

```

```
Parallel = 1
)
Disk diskres10 (
    Partition @primary2 = "/dev/rdisk/c2t50060E8006D43C01d3s2"
    Partition @alternate2 = "/dev/rdisk/c0t50060E8006D43C01d3s2"
)
Disk diskres11 (
    Partition @primary2 = "/dev/rdisk/c2t50060E8005650921d0s2"
    Partition @alternate2 = "/dev/rdisk/c0t50060E8005650921d0s2"
)
Disk diskres12 (
    Partition @primary2 = "/dev/rdisk/c2t50060E8005650921d1s2"
    Partition @alternate2 = "/dev/rdisk/c1t50060E8005650921d1s2"
)
Phantom phres8 (
)
// resource dependency tree
//
// group Storage_group_ldm2_dom2
// {
// Disk diskres10
// Disk diskres11
// Disk diskres12
// Phantom phres8
// }
```



## Support

For technical assistance, visit: [http://www.symantec.com/enterprise/support/assistance\\_care.jsp](http://www.symantec.com/enterprise/support/assistance_care.jsp)

Select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service

## About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at [www.symantec.com](http://www.symantec.com).

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation  
World Headquarters  
20330 Stevens Creek Boulevard  
Cupertino, CA 95014 USA  
+1 (408) 517 8000  
1 (800) 721 3934  
[www.symantec.com](http://www.symantec.com)

Copyright © 2014 Symantec Corporation. All rights reserved. Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

03/14