Lift and Shift Guide - Migrating Workloads from Oracle Solaris 10 (ZFS) SPARC Systems to Oracle Solaris 10 Guest Domains

Part No: E97847-02
February 2020

Lift and Shift Guide - Migrating Workloads from Oracle Solaris 10 (ZFS) SPARC Systems to Oracle Solaris 10 Guest Domains

Part No: E97847-02

Access to Oracle Support

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

Référence: E97847-02

Accès aux services de support Oracle

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

Contents

Using This Documentation

- **Overview** – This document provides instructions for lifting and shifting a SPARC-based Oracle Solaris 10 OS with ZFS root file systems and workloads to a SPARC-based Oracle Solaris 10 Guest Domain in Oracle Solaris 11.3 (or later).
- **Audience** – Experienced Oracle Solaris system administrators
- **Required knowledge** – Experience administering Oracle Solaris computer systems.

Product Documentation Library

Documentation and resources for this product and related products are available at https://docs.oracle.com/cd/E94980_01.

Feedback

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

Understanding the Lift and Shift Process

These topics describe the lift and shift process:

For additional lift and shift guides, visit the Oracle Solaris on SPARC - Lift and Shift Documentation Library at `https://docs.oracle.com/cd/E94980_01/`

Lift and Shift Overview

This document provides instructions for performing a physical to virtual lift and shift. The techniques in this document are specifically aimed at migrating workloads running on Oracle Solaris 10 with ZFS root file systems on SPARC systems to Oracle Solaris 10 Guest Domains in Oracle Solaris 11.3 (or later) running on more modern sun4v hardware.

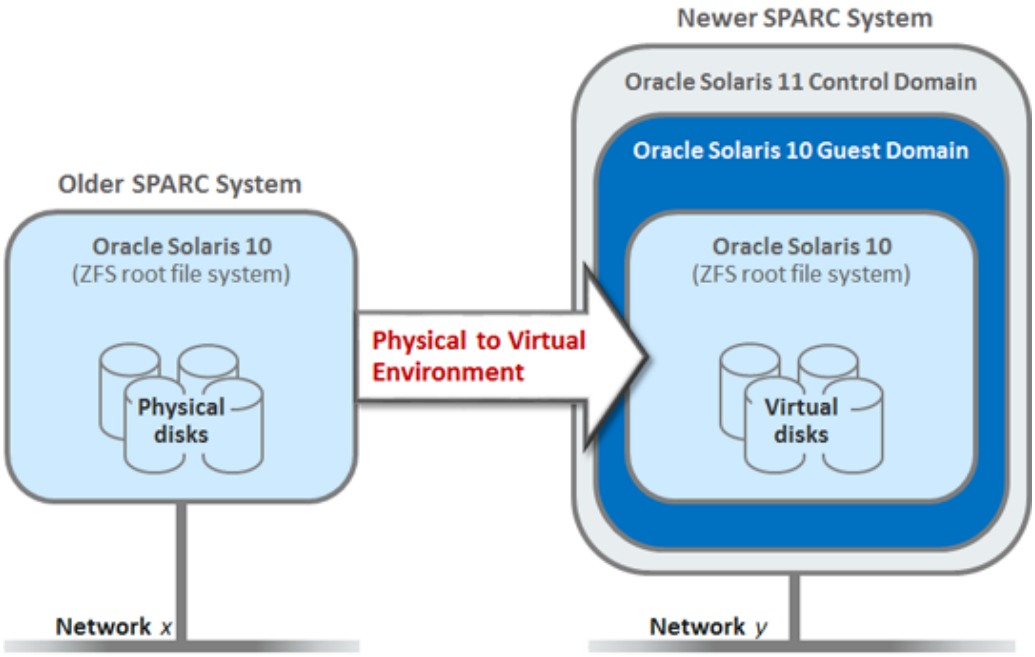

Scope

The instructions in this document are limited to the lift and shift of an Oracle Solaris 10 environment with ZFS root file systems to an Oracle Solaris 11 Environment that is configured with an Oracle Solaris 10 guest domain. The source Oracle Solaris 10 environment can be configured with Solaris zones. In such cases, such as the example in this document, the global zone, non-global zones, and associated applications and data are migrated to the target system. It is also possible to use this guide to migrate a non-virtualized Oracle Solaris 10 environment, as long as the environment is based on ZFS root file systems.

This procedure does not take down the source machine during the archive creation, therefore the source system is available for use. However, to maintain the integrity of the application (Oracle Database and Apache web server in this case), it is a best practice to cleanly shutdown the services.

The performance measurement of the source and target system is not within the scope of this document.

Lift and Shift Process

The lift and shift process involves using a variety of Oracle Solaris utilities to migrate the entire source configuration, including all associated applications and data, to an Oracle Solaris 10 guest domain on the target system. To accomplish this migration, the following activities are performed, although not necessarily in this exact order. The procedures in subsequent chapters order the activities to minimize the amount of time that services are unavailable.

- **Identify and resolve all source system issues** – (Not covered in this guide) Resolving issues before migration simplifies troubleshooting by eliminating the migration as a possible cause of issues. Consider checking log files for issues and applying the latest critical path updates (see Critical Patch Update Bulletins). In some cases, rebooting a system that hasn't been rebooted in a long time can reveal issues.
- **Assess the source system** – Determine the amount of resources used to support the workloads. Gain an understanding of the configuration so that you can replicate the configuration on the target system. Collect information about the CPU, memory, storage, and workload resources.
- **Select the target system** – Select a target system that has sufficient, or additional resources to support the workloads. The target system is usually a more modern system that provides security and performance improvements, and possible cost savings such as lower power, cooling, and support costs.
- **Prepare the source system** – Patches, which provide the necessary lift and shift utilities, are installed on the source system.
- **Prepare the target system** – CPU, memory, and storage resources are configured to accommodate the incoming Oracle Solaris 10 environment.
- **Prepare shared storage** – Shared storage is used to capture the source workloads and then used to restore the workloads on the target system. Using shared storage is central to this lift-and-shift scenario. Storage on the target system is exported as a network file system and mounted on the source system, providing easy access to common storage from the source or target system.
- **Create an Oracle Solaris 10 guest domain on the target system** – A new Oracle Solaris 10 guest domain is created on the target system using an Oracle provided OVM template. The guest domain will host the incoming Oracle Solaris 10 environment. The following illustration shows the commands that are used to prepare the guest domain.

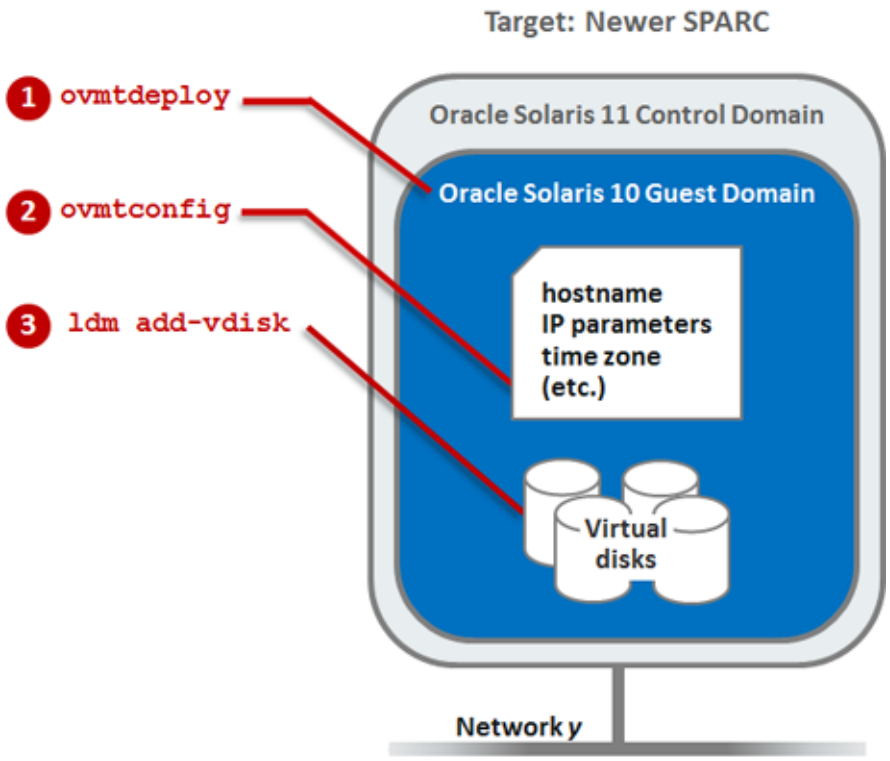

- **Lift** – The `ldmp2vz_collect` command is used to capture the source system's OS ZFS file system, and capture the source's zone configuration (if present) into a flash archive (FLAR). Then, using various utilities, all zones, along with their associated storage content, applications, and application data are lifted from the source and copied to the shared storage.
- **Shift** – The `ldmp2vz_convert` command uses the FLAR that is on the shared storage to restore the source system's OS ZFS file systems to the target guest domain in a new boot environment. The zone configuration and storage is also restored. If the zone storage is included in the root file systems, then that storage is also restored.

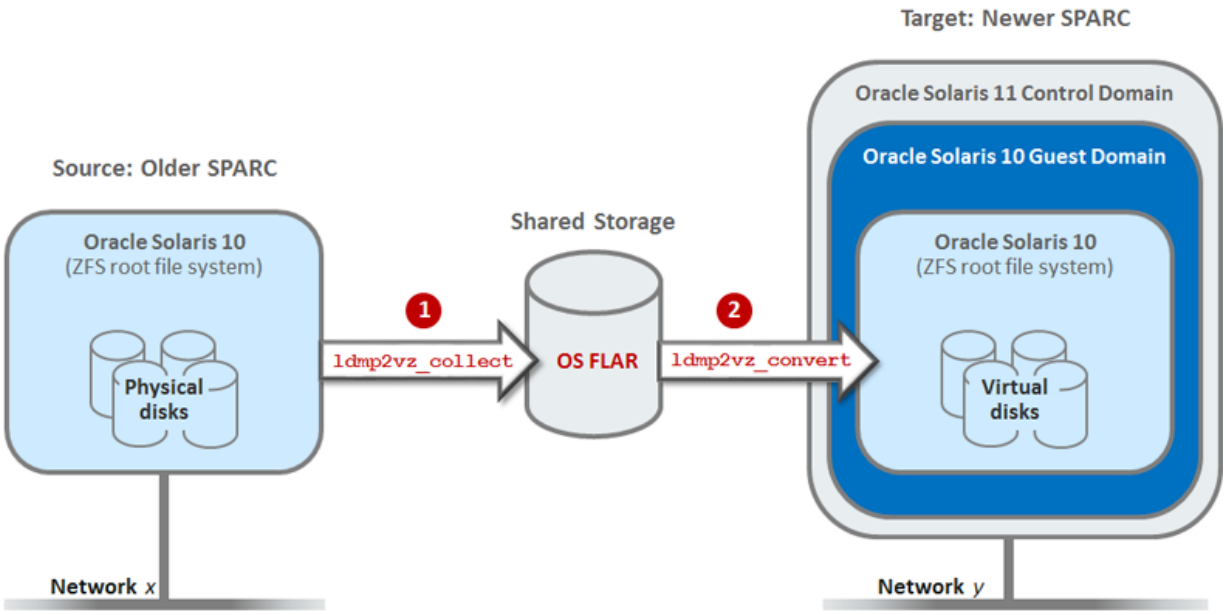

- **Transfer Applications and Data** – Different utilities are used to transfer workload components based on the type of file system where the component resides. In some cases, the output of one command is piped to `pigz` or `gunzip` to improve efficiency through compression and decompression that those commands provide. This diagram shows the different utilities.

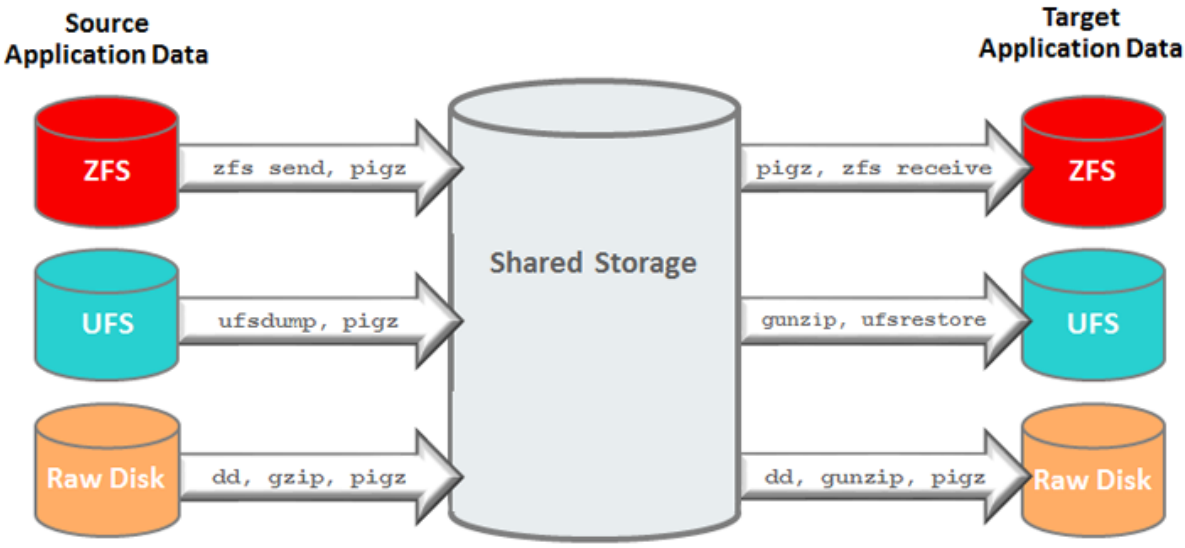

- **Reconfigure** – The shifted environment requires post-migration reconfiguration to function in the new environment. For example, the network configuration is carried over from the source environment, and usually needs to be changed to function in the new environment. If processor sets were used in the source environment, they might need some adjustment when they are restored in the target environment.

Requirements

The lift and shift process described in this document has the following requirements, and includes instructions for satisfying them:

Source System

- Installation of the following patches on the source system (covered in "Prepare the Source System" on page 46):
  - The `SUNWldm` package – Oracle VM Server for SPARC 3.2.
  - The `SUNWldmp2v` package – Oracle VM Server for SPARC 3.2.
  - Patch 119534-33 (or later) – Flash archive patch.

- Patch 151934-06 (or later) – VM Server for SPARC 3.2 ldmd patch, provides lift and shift utilities such as `ldmp2vz_collect`, `ldmp2vz_convert`, and `pigz`.

Target System

- The target control domain must be running Oracle Solaris 11.3 SRU 35 (or later) and LDOM 3.5.0.3.3 (or later).
- Installation of the following patches in the target guest domain (covered in "Patch the Target Guest Domain" on page 43):
  - The latest available Oracle Solaris 10 Recommended patch set (`10_Recommended_CPU_2018-01` at minimum). The patch set provides proactive patching and ensures that the commands in this document run as expected. For more information refer to Critical Patch Updates and MOS knowledge Doc ID 1272947.1 from My Oracle Support (`https://support.oracle.com`).

    Note - This document uses the 10_Recommended_CPU_2018-10 version of this patch set.

  - The `SUNWldm` package – Oracle VM Server for SPARC 3.2.
  - The `SUNWldmp2v` package – Oracle VM Server for SPARC 3.2.
  - Patch 119534-33 (or later) – Flash archive patch.
  - Patch 151934-06 (or later) – VM Server for SPARC 3.2 ldmd patch, provides lift and shift utilities such as `ldmp2vz_collect`, `ldmp2vz_convert`, and `pigz`.
- Plan to download the Oracle VM Template for Oracle VM Server with Oracle Solaris 10.1/13 (Solaris10 OVA) to the target system (covered in "Create a Target Guest Domain" on page 39).

Shared Storage

- On the shared storage, ensure that there is sufficient storage space to temporarily store the image of the source system that is used to perform the lift and shift. The storage must be accessible to the source and target systems. For space requirements and instructions, see "Prepare the Shared Storage" on page 35.
- Plan to download the Oracle Solaris10 (SPARC) ISO (`sol-10-u11-ga-sparc-dvd.iso`) to shared storage (covered in "Prepare the Shared Storage" on page 35).
- Plan to download and unzip all required patches to shared storage (covered in "Prepare the Shared Storage" on page 35). This enables you to download the patches in one location, and patch the source and target systems from that single location.

Example Configuration and Prompts

The lift and shift process in this document is based on a real in-house lift and shift activity. The examples provided in this document are excerpts from that activity. Command line prompts identify on which system the commands are performed.

These systems are used in the examples:

- **Source system** – A SPARC Enterprise M9000 (sun4u) system running the Oracle Solaris 10 OS with two non-global zones:
  - **`dbzone`** – Running Oracle Database 12.1.0.2
  - **`webzone`** – Running Apache web server on a Solaris sparse root zone (zone storage is in the root file system)
- **Target system** – A SPARC S7-2 (sun4v) system, running the Oracle Solaris 11 OS. During the lift and shift procedures, this system is prepared with an Oracle Solaris 10 guest domain that will host the source OS and workloads (the database on the `dbzone` zone and Apache web server on the `webzone` zone).

The following diagrams show the source system configuration and the lift and shift process used in the examples in this guide.

FIGURE 1 Source System Configuration Used in Examples

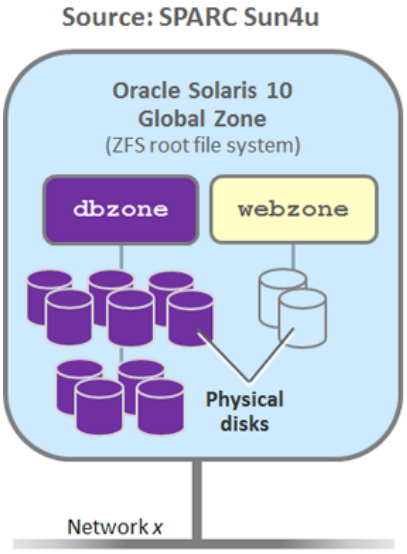

Disk Configuration

Zone	No. LUNs	Data Mgt.	Size	Use
dbzone	2	ZFS	300 GB	DB OS
	2	Raw disk	12 MB	SVM metadb storage
	2	ZFS	150 GB	DB and GI binaries
	2	ASM	200 GB	ASM diskgroups
	2	UFS	250 GB	Redo and archive logs
Global	2	ZFS	600 GB	Global zone OS (also provides storage to the webzone)

FIGURE 2 Example Lift and Shift Process

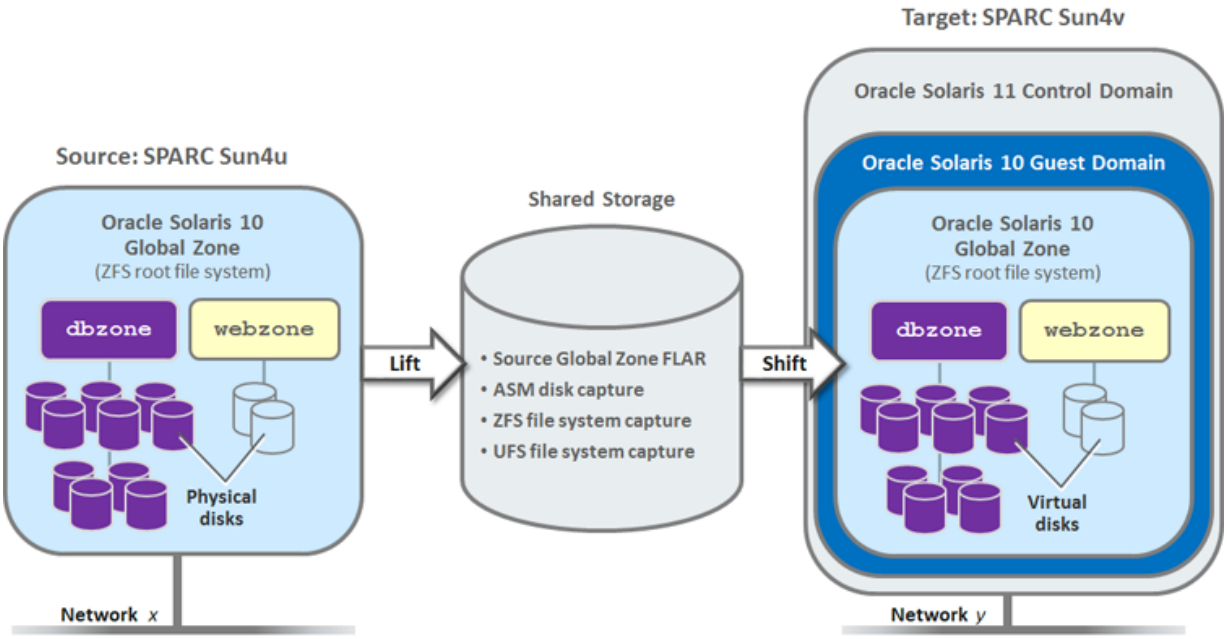

In the screen output examples, the command line prompt indicates on which system (target or source), domain (control or guest), or zone a command is executed. This table lists the prompts.

TABLE 1 System Identification Prompts

System		Host Name in the Prompt
Source System	**Oracle Solaris 10 global zone**	`SourceGlobal`
	dbzone	`SourceDBzone`
	webzone	`SourceWebzone`
Target System	**Oracle Solaris 11 control domain:**	`TargetControlDom`
	Oracle Solaris 10 guest domain:	`TargetGuestDom`
	dbzone:	`TargetDBzone`
	webzone:	`TargetWebzone`

System		Host Name in the Prompt
Other Systems	A client system on the network for testing access to `webzone`.	`client-sys`

Preparing the Lift and Shift Systems

These topics describe how to prepare the source system, the target system, and the shared storage for the lift and shift process:

▼ Review the Source System Configuration

The steps and examples in this section provide commands you can use to determine the configuration of your source system, including details about the zones that you plan to lift and shift. After the lift and shift process, the configuration information can also serve as a reference when you verify the migrated zones on the target system.

As you perform this procedure, take into account the state of your source system and adjust and omit steps as needed.

1. **(Optional) On the source system, start a process that captures the output that is collected in this task.**

   Capturing the commands and output provides a means to refer back to the data that is collected.

   There are a variety of methods to capture output. You can run the script(1M) command to make a record of a terminal session, or use a terminal window with command and output collection capabilities.

   Example:

   ```
   root@SourceGlobal# script /tmp/source_zones_output.txt
   ```

   Note – When you want to stop capturing output, type Ctr-D.

2. List the source system architecture and memory.

```
root@SourceGlobal# uname -a
SunOS SourceGlobal 5.10 Generic_150400-61 sun4u sparc SUNW,SPARC-Enterprise
root@SourceGlobal# prtdiag|grep "Memory size"
Memory size: 163840 Megabytes
```

3. List the swap space configuration.

```
root@SourceGlobal# swap -l
swapfile                 dev    swaplo   blocks    free
/dev/zvol/dsk/rpool/swap 256,1      16   335544304 335544304

root@SourceGlobal # zfs list rpool/swap
NAME        USED  AVAIL  REFER  MOUNTPOINT
rpool/swap  165G   344G   160G  -
```

4. List the dump space configuration.

```
root@SourceGlobal# zfs list rpool/dump
NAME         USED    AVAIL  REFER   MOUNTPOINT
rpool/dump   3.01G   339G   3.00G   -
```

5. List the network information.

```
root@SourceGlobal# ifconfig -a
lo0: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
    inet 127.0.0.1 netmask ff000000
bge0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
    inet 203.0.113.5 netmask ffffff00 broadcast 203.0.113.255
    groupname ipmp0
    ether 0:b:5d:dc:2:3f
bge2: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
    inet 0.0.0.0 netmask ff000000 broadcast 0.255.255.255
    groupname ipmp0
    ether 0:b:5d:dc:2:3f
sppp0: flags=10010008d1<UP,POINTOPOINT,RUNNING,NOARP,MULTICAST,IPv4,FIXEDMTU> mtu 1500
 index 4
    inet 203.0.113.2 --> 203.0.113.1 netmask ffffff00
    ether 0:0:0:0:0:0

root@SourceGlobal# ls -l /etc/hostname*
-rw-r--r--   1 root     root          46 May 16 15:55 /etc/hostname.bge0
-rw-r--r--   1 root     root          16 May 16 15:56 /etc/hostname.bge2

root@SourceGlobal# cat /etc/hostname.bge0
SourceGlobal netmask + broadcast + group ipmp0 up

root@SourceGlobal# cat /etc/hostname.bge2
```

```
group ipmp0 up
```

6. **List the zones.**

```
root@SourceGlobal# zoneadm list -cv
  ID NAME             STATUS     PATH                           BRAND    IP
   0 global           running    /                              native   shared
   1 webzone          running    /rpool/webzone                 native   shared
   2 dbzone           running    /zones/dbzone                  native   shared
```

7. **List the zpool configuration.**

This example displays the zpool configuration information including the zone names, state, and storage information.

```
root@SourceGlobal# zpool status
  pool: dbzone
 state: ONLINE
 scan: resilvered 94K in 0h0m with 0 errors on Fri Apr  6 12:08:40 2018
config:
        NAME                                      STATE     READ WRITE CKSUM
        dbzone                                    ONLINE       0     0     0
          mirror-0                                ONLINE       0     0     0
            c2t600144F0E635D8C700005AC56A460011d0  ONLINE       0     0     0
            c2t600144F0E635D8C700005AC56A6A0012d0  ONLINE       0     0     0

errors: No known data errors

  pool: dbzone_db_binary
 state: ONLINE
 scan: resilvered 94K in 0h0m with 0 errors on Fri Apr  6 12:10:24 2018
config:

        NAME                                      STATE     READ WRITE CKSUM
        dbzone_db_binary                          ONLINE       0     0     0
          mirror-0                                ONLINE       0     0     0
            c2t600144F0E635D8C700005AC56AB30013d0  ONLINE       0     0     0
            c2t600144F0E635D8C700005AC56ADA0014d0  ONLINE       0     0     0

errors: No known data errors

  pool: rpool
 state: ONLINE
 scan: resilvered 11.7G in 0h4m with 0 errors on Wed Apr  4 19:27:39 2018
config:

        NAME          STATE     READ WRITE CKSUM
        rpool         ONLINE       0     0     0
```

```
          mirror-0    ONLINE       0     0     0
            c0t1d0s0  ONLINE       0     0     0
            c1t1d0s0  ONLINE       0     0     0
```

8. Display information about the resource pools.

```
root@SourceGlobal# zonecfg -z dbzone info pool
pool: dbzone-pool
root@SourceGlobal# zonecfg -z webzone info pool
pool: webzone-pool

root@SourceGlobal# poolstat
                             pset
 id pool                 size used load
  1 dbzone-pool            32 0.00 0.05
  0 pool_default           16 0.00 0.04
  2 webzone-pool           32 0.00 0.00
```

9. List the source system processors.

```
root@SourceGlobal# psrinfo -pv
The physical processor has 4 virtual processors (0-3)
  SPARC64-VI (portid 1024 impl 0x6 ver 0x92 clock 2280 MHz)
The physical processor has 4 virtual processors (40-43)
  SPARC64-VI (portid 1064 impl 0x6 ver 0x92 clock 2280 MHz)
The physical processor has 4 virtual processors (80-83)
  SPARC64-VI (portid 1104 impl 0x6 ver 0x92 clock 2280 MHz)
The physical processor has 4 virtual processors (120-123)
  SPARC64-VI (portid 1144 impl 0x6 ver 0x92 clock 2280 MHz)
The physical processor has 8 virtual processors (128-135)
  SPARC64-VII (portid 1152 impl 0x7 ver 0x90 clock 2520 MHz)
The physical processor has 8 virtual processors (136-143)
  SPARC64-VII (portid 1160 impl 0x7 ver 0x90 clock 2520 MHz)
The physical processor has 8 virtual processors (144-151)
  SPARC64-VII (portid 1168 impl 0x7 ver 0x90 clock 2520 MHz)
The physical processor has 8 virtual processors (152-159)
  SPARC64-VII (portid 1176 impl 0x7 ver 0x90 clock 2520 MHz)
The physical processor has 8 virtual processors (160-167)
  SPARC64-VII (portid 1184 impl 0x7 ver 0xa0 clock 2880 MHz)
The physical processor has 8 virtual processors (168-175)
  SPARC64-VII (portid 1192 impl 0x7 ver 0xa0 clock 2880 MHz)
The physical processor has 8 virtual processors (176-183)
  SPARC64-VII (portid 1200 impl 0x7 ver 0xa0 clock 2880 MHz)
The physical processor has 8 virtual processors (184-191)
  SPARC64-VII (portid 1208 impl 0x7 ver 0xa0 clock 2880 MHz)
```

10. List the processors assigned to `dbzone`.

The dbzone is assigned the last 4 physical processors (the last 32 virtual processors, CPU IDs 160 to 191).

```
root@SourceGlobal# zlogin dbzone psrinfo -pv
The physical processor has 8 virtual processors (160-167)
  SPARC64-VII (portid 1184 impl 0x7 ver 0xa0 clock 2880 MHz)
The physical processor has 8 virtual processors (168-175)
  SPARC64-VII (portid 1192 impl 0x7 ver 0xa0 clock 2880 MHz)
The physical processor has 8 virtual processors (176-183)
  SPARC64-VII (portid 1200 impl 0x7 ver 0xa0 clock 2880 MHz)
The physical processor has 8 virtual processors (184-191)
  SPARC64-VII (portid 1208 impl 0x7 ver 0xa0 clock 2880 MHz)
```

11. List the processors assigned to `webzone`.

The webzone is assigned 4 physical processors (32 virtual processors, CPU IDs 128 to 159). These processors are the set of processors before the dbzone processors.

```
root@SourceGlobal# zlogin webzone psrinfo -pv
The physical processor has 8 virtual processors (128-135)
  SPARC64-VII (portid 1152 impl 0x7 ver 0x90 clock 2520 MHz)
The physical processor has 8 virtual processors (136-143)
  SPARC64-VII (portid 1160 impl 0x7 ver 0x90 clock 2520 MHz)
The physical processor has 8 virtual processors (144-151)
  SPARC64-VII (portid 1168 impl 0x7 ver 0x90 clock 2520 MHz)
The physical processor has 8 virtual processors (152-159)
  SPARC64-VII (portid 1176 impl 0x7 ver 0x90 clock 2520 MHz)
```

12. Identify the source system disks and what is stored on them.

a. Use the `format` utility to list the disks.

```
root@SourceGlobal# format

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
          /pci@0,600000/pci@0/scsi@1/sd@0,0
       1. c0t1d0 <SEAGATE-ST960004SSUN600G-0115 cyl 64986 alt 2 hd 27 sec 668>
          /pci@0,600000/pci@0/scsi@1/sd@1,0
       2. c1t0d0 <SUN600G cyl 64986 alt 2 hd 27 sec 668>
          /pci@24,600000/pci@0/scsi@1/sd@0,0
       3. c1t1d0 <SUN600G cyl 64986 alt 2 hd 27 sec 668>
          /pci@24,600000/pci@0/scsi@1/sd@1,0
       4. c2t600144F0E635D8C700005AC7C46F0017d0 <SUN-ZFSStorage7420-1.0 cyl 8124 alt
 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f0e635d8c700005ac7c46f0017
       5. c2t600144F0E635D8C700005AC7C4920018d0 <SUN-ZFSStorage7420-1.0 cyl 8124 alt
 2 hd 254 sec 254>
```

```
          /scsi_vhci/ssd@g600144f0e635d8c700005ac7c4920018
       6. c2t600144F0E635D8C700005AC56A6A0012d0 <SUN-ZFS Storage 7420-1.0-300.00GB>
          /scsi_vhci/ssd@g600144f0e635d8c700005ac56a6a0012
       7. c2t600144F0E635D8C700005AC56A460011d0 <SUN-ZFS Storage 7420-1.0-300.00GB>
          /scsi_vhci/ssd@g600144f0e635d8c700005ac56a460011
       8. c2t600144F0E635D8C700005AC56AB30013d0 <SUN-ZFS Storage 7420-1.0-150.00GB>
          /scsi_vhci/ssd@g600144f0e635d8c700005ac56ab30013
       9. c2t600144F0E635D8C700005AC56ADA0014d0 <SUN-ZFS Storage 7420-1.0-150.00GB>
          /scsi_vhci/ssd@g600144f0e635d8c700005ac56ada0014
      10. c2t600144F0E635D8C700005AC56B2E0016d0 <SUN-ZFS Storage 7420-1.0-200.00GB>
          /scsi_vhci/ssd@g600144f0e635d8c700005ac56b2e0016
      11. c2t600144F0E635D8C700005AC56B080015d0 <SUN-ZFS Storage 7420-1.0-200.00GB>
          /scsi_vhci/ssd@g600144f0e635d8c700005ac56b080015
CTRL-C
```

b. Identify the amount of storage used by the ZFS file systems.

The values in the `ALLOC` column are used to calculate the amount of storage that is needed in shared storage and on the target system.

```
root@SourceGlobal #zpool list
NAME              SIZE  ALLOC  FREE  CAP  HEALTH  ALTROOT
dbzone            298G  16.4G  282G   5%  ONLINE  -
dbzone_db_binary  149G  22.2G  127G  14%  ONLINE  -
rpool             556G   199G  357G  35%  ONLINE  -
```

c. Identify the amount of storage used by the UFS file systems.

In this example, the dbzone logs are on UFS file systems on disks 4 and 5.

The values in the `used` column are used to calculate the amount of storage that is needed in shared storage and on the target system.

```
root@SourceGlobal# zlogin dbzone df -h -F ufs
Filesystem           size   used    avail   capacity    Mounted on
/logs/archivelogs    197G    13G    182G    7%          /logs/archivelogs
/logs/redologs       9.8G   1.2G    8.6G    13%         /logs/redologs
```

d. Identify the amount of storage used by the ASM raw disks.

In this example, the ASM raw disks are on slice 0 of disks 10 and 11 (as numbered by the `format` utility).

The capacities in GB can be extracted from the disk partition table using the following scripts.

For this example, each ASM raw disk is approximately 199 GB.

```
root@SourceGlobal# echo $((`prtvtoc /dev/rdsk/
c2t600144F0E635D8C700005AC56B2E0016d0s2 | grep " 0 " | awk '{ print
 $5 }'`*512/1024/1024/1024))
199
root@SourceGlobal# echo $((`prtvtoc /dev/rdsk/
c2t600144F0E635D8C700005AC56B080015d0s2 | grep " 0 " | awk '{ print
 $5 }'`*512/1024/1024/1024))
199
```

e. Determine the total amount of storage used.

The disk capacity information is used to configure the target system storage that is needed for the migration. See "Prepare the Target System" on page 29.

The total amount of used storage is used to calculate the amount of space needed in the shared storage in "Prepare the Shared Storage" on page 35.

This table lists the disk calculation for this example.

`format` Disk No.	Disk Capacity (GB)	Data Mgt.	Used Storage	Contents
0, 2	2 x 146 = 292	raw disk	2 x 12 MB	SVM metadb **Note -** SVM metadb data is not copied to the shared storage and is not included in the total used storage value.
1, 3	2 x 600 = 1200	ZFS	199 GB	Global zone OS (includes storage for the `webzone`)
4, 5	2 x 250 = 500	UFS	13 GB 1.2 GB	`dbzone` redo and archive logs
6, 7	2 x 300 = 600	ZFS	16.4 GB	`dbzone` OS
8, 9	2 x 150 = 300	ZFS	22.2 GB	DB and GI binaries
10, 11	2 x 200 = 400	ASM	398 GB	ASM diskgroups
Totals	**3292 GB**		**649.8 GB**	

13. List the SVM metadevice state database (`metadb`) configuration.

In this example, two internal disks (`c0t0d0s4` and `c1t0d0s4`) are used for redundancy of the metadevice state database.

```
root@SourceGlobal# metadb
        flags           first blk       block count
     a m  pc luo        16              8192            /dev/dsk/c0t0d0s4
     a    pc luo        8208            8192            /dev/dsk/c0t0d0s4
     a    pc luo        16400           8192            /dev/dsk/c0t0d0s4
```

```
    a    pc luo        16              8192            /dev/dsk/c1t0d0s4
    a    pc luo        8208            8192            /dev/dsk/c1t0d0s4
    a    pc luo        16400           8192            /dev/dsk/c1t0d0s4
```

14. Display the list of active metadevices.

The file systems displayed in the previous command are carved out of SVM meta devices (which are in turn based on the disks mentioned above) and can be displayed with the `metastat` command.

```
root@SourceGlobal# metastat -p
d30 -p d0 -o 21036096 -b 419430400
d0 -m d11 d12 1
d11 1 1 /dev/dsk/c2t600144F0E635D8C700005AC7C46F0017d0s0
d12 1 1 /dev/dsk/c2t600144F0E635D8C700005AC7C4920018d0s0
d20 -p d0 -o 64544 -b 20971520
```

15. On `dbzone`, list the architecture, network parameters, and file system information.

```
root@SourceDBzone# uname -a
SunOS SourceDBzone 5.10 Generic_150400-61 sun4u sparc SUNW,SPARC-Enterprise

root@SourceDBzone# ifconfig -a
lo0:1: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
        inet 127.0.0.1 netmask ff000000
bge0:2: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
        inet 203.0.113.5 netmask ffffff00 broadcast 203.0.113.255

root@SourceDBzone# netstat -rn
Routing Table: IPv4
  Destination           Gateway         Flags  Ref  Use       Interface
-------------------- -------------------- ----- ----- --------- ---------
default              203.0.113.1        UG     1    186588
203.0.113.0          203.0.113.5        U      1         2    bge0:2
224.0.0.0            203.0.113.5        U      1         0    bge0:2
127.0.0.1            127.0.0.1          UH     7        54    lo0:1

root@SourceDBzone# zfs list
NAME                     USED  AVAIL  REFER  MOUNTPOINT
dbzone_db_binary       22.2G   124G  22.2G  /u01
```

16. Display the database disk information for `dbzone`.

Log into the `dbzone` as an `oracle` grid user and enter `asmcmd` to start the ASMCMD command line.

Note that the ownership of the device slices (s0) for ASM is `oracle:dba` and relates to the Oracle installation in the non-global zone.

```
ASMCMD> lsdg
State    Type    Rebal  Sector  Block       AU  Total_MB  Free_MB  Req_mir_free_MB
 Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED  NORMAL  N         512   4096  16777216    409568   356960                0
     178480              0             N  DATA/

ASMCMD> lsdsk -p -G DATA
Group_Num  Disk_Num      Incarn  Mount_Stat  Header_Stat  Mode_Stat  State   Path
        1         0  4239149010  CACHED      MEMBER       ONLINE     NORMAL  /dev/rdsk/
c2t600144F0E635D8C700005AC56B080015d0s0
        1         1  4239149009  CACHED      MEMBER       ONLINE     NORMAL  /dev/rdsk/
c2t600144F0E635D8C700005AC56B2E0016d0s0
```

17. On `webzone`, list the architecture, network parameters, and file system information.

```
root@SourceWebzone# uname -a
SunOS SourceWebzone 5.10 Generic_150400-61 sun4u sparc SUNW,SPARC-Enterprise

root@SourceWebzone# ifconfig -a
lo0:3: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
        inet 127.0.0.1 netmask ff000000
bge0:4: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
        inet 203.0.113.191 netmask ffffff00 broadcast 203.0.113.255
root@SourceWebzone# netstat -rn
Routing Table: IPv4
  Destination           Gateway           Flags  Ref     Use     Interface
-------------------- -------------------- ----- ----- ---------- ---------
default              203.0.113.1         UG       1     186609
203.0.113.0          203.0.113.191       U        1          2    bge0:3
224.0.0.0            203.0.113.191       U        1          0    bge0:3
127.0.0.1            127.0.0.1           UH       4         89    lo0:2

root@SourceWebzone# zfs list
no datasets available
```

▼ Prepare the Target System

Use this procedure to ensure that the target system is configured to provide CPU, memory, and storage resources for the incoming source environment.

- **CPU Resources** – During the shift, you can assign any amount of CPU resources to the guest domain that are appropriate for the guest domain's workload. However, prior to the lift and shift, you must ensure that those CPU resources are available as described in this procedure.

If you are uncertain about the CPU utilization of the guest domain's workload on the target system, then the target system should provide at minimum the same available CPU and memory resources as on the source system. This conservative approach helps maintain the same or better performance level of the workload after the migration. On the other hand, if the CPU utilization is estimated to be significantly lower for the guest domain on the target system, for example, if the target system has faster CPUs, then the target system can provide fewer CPU resources to the guest domain. In some cases, using fewer CPU cores reduces software licensing costs.

- **Memory Resources** – By default, the target guest domain is allocated the same amount of memory that exists on the source system. Ensure that there is at least the same amount of memory available on the target system as described in this procedure.
- **Storage Resources** – The number of available virtual disks and sizes on the target must match the source system's physical disks, with one exception. The disks that store the SVM metadb data can be different. This procedure describes how to configure the target storage.
- **Software Version Requirements** – The target control domain must be running Oracle Solaris 11.3 SRU 35 (or later) and LDOM 3.5.0.3.3 (or later). The first step in this procedure describes how to identify these software versions.

1. Verify that the target system is running the required software versions.

If your target system does not have the minimum required software versions, update the target system before continuing. Refer to the applicable Oracle Solaris documentation for updating instructions:

- Oracle Solaris 11.3 Information Library
- Oracle Solaris 11.4 Information Library

```
root@TargetControlDom# uname -a
SunOS TargetControlDom 5.11 11.3 sun4v sparc sun4v

root@TargetControlDom# pkg list entire ldomsmanager
NAME                          (PUBLISHER)                    VERSION IFO
entire                         0.5.11-0.175.3.35.0.6.0       i--
system/ldoms/ldomsmanager      3.5.0.3.3-0.175.3.35.0.4.0    i--

root@TargetControlDom# ldm -V

Logical Domains Manager (v 3.5.0.3.3)
        Hypervisor control protocol v 1.12
        Using Hypervisor MD v 1.4

System PROM:
        Hostconfig v. 1.8.3.a @(#)Hostconfig 1.8.3.a 2016/09/16 14:15
        Hypervisor v. 1.17.3.a @(#)Hypervisor 1.17.3.a 2016/09/16 13:38
        OpenBoot v. 4.40.3 @(#)OpenBoot 4.40.3 2016/08/17 12:17
```

2. Check the target system's storage capacity.

In this example, the LUNS on the target system that are intended for the guest domain are provisioned with the exact same capacity as in the source system (see "Review the Source System Configuration" on page 21).

This table lists the disk configuration for this example.

`format` Disk No.	Size (GB)	Contents
11, 16	2 x 1	SVM metadb storage
14, 15	2 x 250	`dbzone` redo and archive logs
5, 8	2 x 300	`dbzone` OS
12, 13	2 x 150	DB and GI binaries
9, 10	2 x 200	ASM diskgroups
6, 7	2 x 600	Global zone OS (includes storage for the `webzone`)
	2 x 1501 = 3002	TOTAL Disk Space

a. Obtain disk names.

In this example, disk names are displayed as `c0t5000CCA08041FCF4d0`, `c0t600144F09F2C0BFD00005BE4A9A90005d0`, and so on.

```
root@TargetControlDom# echo | format
AVAILABLE DISK SELECTIONS:

       0. c0t5000CCA08041FCF4d0 <HGST-H101812SFSUN1.2T-A990-1.09TB>
          /scsi_vhci/disk@g5000cca08041fcf4
          /dev/chassis/SYS/HDD0/disk
       1. c0t5000CCA080409B54d0 <HGST-H101812SFSUN1.2T-A990-1.09TB>
          /scsi_vhci/disk@g5000cca080409b54
          /dev/chassis/SYS/HDD1/disk
       2. c0t5000CCA080409D24d0 <HGST-H101812SFSUN1.2T-A990-1.09TB>
          /scsi_vhci/disk@g5000cca080409d24
          /dev/chassis/SYS/HDD2/disk
       3. c0t5000CCA080409F14d0 <HGST-H101812SFSUN1.2T-A990-1.09TB>
          /scsi_vhci/disk@g5000cca080409f14
          /dev/chassis/SYS/HDD3/disk
       4. c1t0d0 <MICRON-eUSB DISK-1112 cyl 246 alt 0 hd 255 sec 63>
          /pci@300/pci@1/pci@0/pci@2/usb@0/storage@1/disk@0,0
          /dev/chassis/SYS/MB/EUSB_DISK/disk
       5. c0t600144F09F2C0BFD00005BE4A9A90005d0 <SUN-ZFS Storage 7355-1.0 cyl 9749
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4a9a90005
```

```
       6. c0t600144F09F2C0BFD00005BE4A90F0004d0 <SUN-ZFS Storage 7355-1.0 cyl 19501
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4a90f0004
       7. c0t600144F09F2C0BFD00005BE4A8500003d0 <SUN-ZFS Storage 7355-1.0 cyl 19501
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4a8500003
       8. c0t600144F09F2C0BFD00005BE4BF670006d0 <SUN-ZFS Storage 7355-1.0 cyl 9749
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4bf670006
       9. c0t600144F09F2C0BFD00005BE4BFC90007d0 <SUN-ZFS Storage 7355-1.0 cyl 6499
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4bfc90007
      10. c0t600144F09F2C0BFD00005BE4BFE60008d0 <SUN-ZFS Storage 7355-1.0 cyl 6499
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4bfe60008
      11. c0t600144F09F2C0BFD00005BE4C3EC000Bd0 <SUN-ZFS Storage 7355-1.0 cyl 8190
 alt 2 hd 8 sec 32>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4c3ec000b
      12. c0t600144F09F2C0BFD00005BE4C4BB000Dd0 <SUN-ZFS Storage 7355-1.0 cyl 4873
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4c4bb000d
      13. c0t600144F09F2C0BFD00005BE4C4E4000Ed0 <SUN-ZFS Storage 7355-1.0 cyl 4873
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4c4e4000e
      14. c0t600144F09F2C0BFD00005BE4C06B000Ad0 <SUN-ZFS Storage 7355-1.0 cyl 8124
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4c06b000a
      15. c0t600144F09F2C0BFD00005BE4C0410009d0 <SUN-ZFS Storage 7355-1.0 cyl 8124
 alt 2 hd 254 sec 254>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4c0410009
      16. c0t600144F09F2C0BFD00005BE4C424000Cd0 <SUN-ZFS Storage 7355-1.0 cyl 8190
 alt 2 hd 8 sec 32>
          /scsi_vhci/ssd@g600144f09f2c0bfd00005be4c424000c
Specify disk (enter its number): Specify disk (enter its number):
```

b. Display the disk capacities.

The disk topology and capacities must match the disk topology and capacities as the source system.

You can use the `iostat -En` command with each disk name that was provided in the previous step. For example:

`iostat -En` *disk_name*

Repeat the `iostat` command for each disk.

Note that the sizes shown represent the raw whole disk capacity including a reserved area. The actual usable capacity is less.

This example shows that the target system is configured to provide the exact same virtual disks and capacities that the guest domain had on the source system.

```
root@TargetControlDom# iostat -En c0t600144F09F2C0BFD00005BE4C424000Cd0 | grep -i
 size
Size: 1.07GB <1073741824 bytes>
```

3. Configure CPU and memory resources to support the incoming guest domain.

a. Ensure that the path of the `ovmtdeploy` command is in your `PATH` variable.

For example:

```
root@TargetControlDom# export PATH=$PATH:/opt/ovmtutils/bin
```

b. List the target system's CPU and memory resources.

If your system has multiple domains, add the vCPU and memory resources for all domains to determine the total allocated resources.

In this example, the target system is set to factory defaults, and all of the resources are currently assigned to the control domain. To make resources available for the guest domain, some resources must be removed from the control domain.

```
root@TargetControlDom# ldm ls
NAME       STATE    FLAGS   CONS     VCPU   MEMORY    UTIL   NORM   UPTIME
primary    active   -n-c--  UART     128    260352M   0.3%   0.3%   1d
```

Note that the `ldm ls` command only displays resources that are allocated to domains. If there are unallocated resources, they are not displayed. If you need to identify unallocated resources, run these commands.

- List the number of unallocated cores:

  # **`ldm list-devices -p core | grep cid | wc -l`**
- List the unallocated memory (add the values displayed under the size column):

  # **`ldm list-devices memory`**

c. Set aside CPU and memory resources for the guest domain.

In this example, the target control domain's resources are reduced to free up resources for the incoming guest domain, while maintaining an optimal amount of resources to continue providing services.

```
root@TargetControlDom# ldm set-core 2 primary
root@TargetControlDom# ldm set-mem 32G primary
```

4. **Save the new logical domain resource configuration.**

   This ensures that the configuration is saved and used after a power cycle. If it isn't saved, the configuration is lost after a power cycle.

   a. **Add a logical domain configuration to the service processor.**

   ```
   root@TargetControlDom# ldm add-spconfig initial
   ```

   b. **Check the current logical domain configuration on the service processor.**

   In this example, the output shows that the newly created `initial` configuration is now active.

   ```
   root@TargetControlDom# ldm ls-spconfig

   factory-default
   initial [current]
   ```

   c. **Reboot the target system.**

   ```
   root@TargetControlDom# shutdown -i 6
   ```

   d. **After the system boots, log in and check the SP configuration.**

   Ensure that the new configuration is the current configuration.

   ```
   root@TargetControlDom# ldm ls-spconfig

   factory-default
   initial [current]
   ```

5. **Verify the CPU and memory configuration.**

   In this example, the output shows that the control domain has reduced resources. Unallocated resources are now available to be allocated to the incoming guest domain.

   ```
   root@TargetControlDom# ldm ls

   NAME             STATE      FLAGS   CONS    VCPU  MEMORY   UTIL  NORM  UPTIME
   primary          active     -n-cv-  UART    16    32G      0.7%  0.7%  1m
   ```

6. **Verify the network redundancy that will support the incoming workloads.**

On the source system, network redundancy was configured using IPMP. On the target system, link aggregation is used to follow best practices when deploying logical domains that use networking from one service domain.

In this example, network redundancy is provided by `aggr0` in the control domain.

```
root@TargetControlDom# dladm show-aggr
LINK              MODE  POLICY   ADDRPOLICY           LACPACTIVITY LACPTIMER
aggr0             trunk L4       auto                 active       short

root@TargetControlDom# dladm show-aggr -L
LINK                PORT         AGGREGATABLE SYNC COLL DIST DEFAULTED EXPIRED
aggr0               net0         yes          yes  yes  yes  no        no
--                  net2         yes          yes  yes  yes  no        no
```

▼ Prepare the Shared Storage

In this procedure, a network file system (NFS) is exported from the target system and mounted by the source system. During the lift process, the file system archive file is created on the shared storage where it is accessible to both the source and target systems.

Requirements

- **Shared storage** – The location must be accessible to the source system and the target system.
- **Sufficient storage space** – The shared storage must have enough available storage space to temporarily store these components:
  - Operating system FLAR
  - Application data (ASM)
  - Application specific logs (redo and archive)
  - Application binaries (ZFS streams)
  - Oracle Solaris10 DVD (ISO required to change from sun4u to sun4v.)
  - OVM template (when unzipped, is approximately 2.6 GB)
  - Unzipped patches: 10_Recommended_CPU_2018-10 and 151934-06 (approximately 14 GB)

For more information about networked file systems, refer to these resources:

- Oracle Solaris 10, Managing File Systems at: https://docs.oracle.com/cd/E23823_01/html/816-4555/rfsintro-2.html

- Oracle Solaris 11, Managing Network File Systems at: https://docs.oracle.com/cd/E23824_01/html/821-1454/rfsintro-2.html

1. Ensure that there is enough available storage capacity on the planned shared storage.

Use the storage information gathered in "Review the Source System Configuration" on page 21.

a. Calculate the space needed to store the capture of source system.

Take the source system's total used storage (obtained in "Review the Source System Configuration" on page 21) and divide by two. This calculation takes into account that the various data captured items will be compressed in the shared storage, yet is conservative enough to ensure adequate storage space for all the items.

Item	Calculation	Total
Source system's total used storage, divided by 2 (To account for a conservative compression ratio)	649.8 ÷ 2	325 GB
S10u11 DVD ISO		2.2 GB
S10u11 OVM template		2.6 GB
Patches (both zipped and unzipped)		14 GB
Sum of above (rounded)		344 GB
Add approximately 30% (To provide additional headroom for the creation of temporary files, uncompressing files, and so on.)	344 + 100	**544 GB**

For this example scenario, **544 GB** of available space is required on the shared storage.

b. Identify available storage in the target system control domain.

The command output shows that there is approximately 1.09 TB available storage capacity, which is sufficient space to store the compressed images of the source system.

```
root@TargetControlDom# zpool list vpool
NAME    SIZE    ALLOC   FREE    CAP   DEDUP   HEALTH   ALTROOT
vpool   1.09T   24M     1006G   1%    1.00x   ONLINE   -
```

c. (If needed) Add storage space to accommodate the storage space requirements.

2. **On the target control domain, create and export a file system.**

```
root@TargetControlDom# zfs create -o mountpoint=/ovas1 vpool/ovas1
root@TargetControlDom# zfs set share.nfs=on vpool/ovas1
root@TargetControlDom# zfs set share.nfs.root=\* vpool/ovas1
root@TargetControlDom# zfs set share.nfs.rw=\* vpool/ovas1
root@TargetControlDom# exportfs
vpool_ovas1      /ovas1   sec=sys,root,rw
```

3. **Create a directory on shared storage that will be used for downloads.**

```
root@TargetControlDom# mkdir /ovas1/Downloads
```

4. **As superuser on the source system, create a mount point and mount the shared storage.**

   In this example, `TargetControlDom:/ovas1` is an NFS share exported from the target system.

```
root@SourceGlobal# mkdir /ovas1
root@SourceGlobal# mount -F nfs TargetControlDom:/ovas1 /ovas1
```

5. **Download the Oracle Solaris 10 (SPARC) ISO image to the shared storage.**

   The ISO image is called `sol-10-u11-ga-sparc-dvd.iso` and available from https://www.oracle.com/solaris/solaris10/downloads/solaris10-get-jsp-downloads.html

   In this example, all downloaded items, including the ISO image, are downloaded to the `/ovas1/Downloads` directory on the shared storage.

6. **Download and unzip the Template for Oracle VM Server with Oracle Solaris 10 1/13 on the target system.**

   You can access the template by signing in to https://edelivery.oracle.com, and selecting these choices: All categories → Oracle VM Template for Oracle VM Server with Oracle Solaris → REL:Oracle VM Template for Oracle VM Server with Oracle Solaris 10.1.13 → Accept license agreement → Download V37543-01.zip

```
root@TargetControlDom# cd /ovas1/Downloads
root@TargetControlDom# unzip ./V37543-01.zip
```

7. **Capture the source system zones configuration and copy the configuration files to the shared storage.**

   The source system zone configurations are copied for potential future reference. The configuration files are not used for the migration.

```
root@SourceGlobal# zonecfg -z dbzone export > /ovas1/dbzone.cfg
root@SourceGlobal# zonecfg -z webzone export > /ovas1/webzone.cfg
```

8. From the target control domain, download and unzip the required patches on the shared storage.

The patches are later used to patch the source system and target guest and control domains. By downloading the patches to the shared storage, the patches are downloaded once to a single location, but used several times during the lift and shift process.

a. Download the required patches.

The patches and README files are available for download from My Oracle Support (`https://support.oracle.com`).

Download these patch zip files to the shared storage. This example downloads the patches to the `/ovas1/Downloads` directory.

Ensure that you select the SPARC 64 bit versions of the patches.

This table lists the procedures in this document that require patch installation.

Patch ID	Description	Where Patches are Installed
`20997186`	Provides the Oracle VM Server for SPARC 3.2 `SUNWldm` and `SUNWldmp2v` packages in a zipped form. Some of the lift and shift utilities require the installation of these two packages. **Note -** `SUNWldm` and `SUNWldmp2v` must be installed before installing patch 151934.	`20997186` is unzipped on shared storage (next step). From shared storage, the `SUNWldm` and `SUNWldmp2v` packages are installed in these systems: ■ Source global zone in "Prepare the Source System" on page 46. ■ Target guest domain in "Patch the Target Guest Domain" on page 43
`119534-33` (or later)	Flash archive patch.	■ Source global zone in "Prepare the Source System" on page 46. ■ Target guest domain in "Create a Target Guest Domain" on page 39
`151934-06` (or later)	The Oracle VM Server for SPARC 3.2 patch provides lift and shift utilities such as `ldmp2vz_collect`, `ldmp2vz_convert`, and `pigz`	■ Source global zone in "Prepare the Source System" on page 46. ■ Target guest domain in "Create a Target Guest Domain" on page 39
`10_Recommended_CPU_2018-10` (or later)	Provides proactive patching and ensures that the commands in this document run as expected. For information on locating recommended OS patch sets on MOS, refer to the MOS knowledge article with Doc ID 1272947.1.	■ Target guest domain in "Patch the Target Guest Domain" on page 43 and again in "Install the Oracle Solaris 10 Recommended Patch Set" on page 60.

Note - The `SUNWldm`, `SUNWldmp2v`, `119534-33`, and `151934-06` patches do not modify the OS. They only install the utilities used for this lift and shift process.

b. Unzip the patch zip files.

```
root@TargetControlDom# cd ovas1/Downloads
root@TargetControlDom# unzip ./p20997186_32000_SOLARIS64.zip
root@TargetControlDom# unzip ./119534-33.zip
root@TargetControlDom# unzip ./151934-06.zip
root@TargetControlDom# unzip ./10_Recommended_CPU_2018-10.zip
```

▼ Create a Target Guest Domain

In this example, the target system is an Oracle SPARC S7-2 running the Oracle Solaris 11 OS. This procedure describes how to create an Oracle Solaris 10 guest domain that will host the migrated zones.

As you perform this procedure, take into account the state of your target system and adjust or omit steps as needed.

1. On the target control domain, save a new logical domain resource configuration.

This ensures that the configuration is saved and used after a power cycle. If it isn't saved, the configuration is lost after a power cycle.

a. Add a logical domain configuration to the service processor.

```
root@TargetControlDom# ldm add-spconfig solaris10config
```

b. Check that the newly created logical domain configuration is active.

In this example, the output shows that the newly created `solaris10config` configuration is active.

```
root@TargetControlDom# ldm ls-spconfig

factory-default
initial
solaris10config [current]
```

2. On the target system, create an Oracle Solaris 10 guest domain.

The guest domain will be deployed with the source's Oracle Solaris 10 OS and host the zones.

This example uses the `ovmtdeploy` command. Alternatively, you can create the guest domain from an Oracle Solaris DVD or use Jumpstart.

Syntax:

`/opt/ovmtutils/bin/ovmtdeploy -d` *domain-name* `-s -c` *vcpu* `-m` *size* `-t` *vsw* `-e` *nic* `-E` *vnet* `-v` *dev template_name*

Where:

- *-d domain-name* – The name of the guest domain being created
- `-s` – Prevents the new guest domain from automatically starting after its creation. This enables you to reconfigure the guest domain before it starts.
- `-c` *vcpu* – Specifies the number of virtual CPUs to assign (overrides the setting for the number of virtual CPUs specified in the template file). The number of vCPUs must not exceed the number of vCPUs that are available in the target system and intended for this guest domain.
- `-m` *size* – Overrides the memory size specified in the template file (.ovf). The number of overriding sizes must match the number of logical domains. *size* is the amount of memory plus a unit. The unit is `m` for megabytes or `g` for gigabytes.
- `-t` *vsw* – The name of the virtual switch
- `-e` *nic* – Specifies the network adapter to use for the virtual network switches that are required for this template.
- `-E` *vnet* – The name of the virtual network
- `-v` *dev* – Specifies a comma-separated list of target devices, directories, or both (depends on what is specified for the *template_name*).
- *template_name* – The path name to the OVA template file. This file must be an OVF 1.0 archive that contains disk images formatted for use by SPARC systems, downloaded from `https://support.oracle.com` or `https://edelivery.oracle.com`

This single command line creates an Oracle Solaris 10 guest domain:

```
root@TargetControlDom# /opt/ovmtutils/bin/ovmtdeploy -d solaris10 -s -c 80 -m 128g -t
 primary-vsw0 -e aggr0 -E vnet0 -v /dev/rdsk/c0t600144F09F2C0BFD00005BE4A8500003d0s2
  /ovas1/Downloads/OVM_S10U11_SPARC.ova
```

3. Verify the existence of the new target guest domain.

This example shows that the `solaris10` guest domain is created.

```
root@TargetControlDom# ldm ls
NAME             STATE      FLAGS     CONS       VCPU    MEMORY    UTIL   NORM   UPTIME
```

```
primary          active     -n-cv-    UART       16      32G       1.2%   1.2%   6d 14h
 11m
solaris10        bound      -----     5000       80      128G
```

4. Configure the guest domain.

The target system is on a separate network (192.0.2.0) than the Source environment (203.0.113.0). The transition of the subnet from source to target can be performed manually using the `sys-config` command, or by using the `ovmtconfig` utility. These steps use the `ovmtconfig` utility.

a. Create a properties file.

The `ovmtconfig` utility requires a configuration properties file that provides the configuration parameters.

Copy and modify the `/opt/ovmtutils/share/props/solaris.properties` file:

```
root@TargetGuestDom# cp -ip /opt/ovmtutils/share/props/solaris.properties /ovas1/
solaris10.props
```

In this example, the `solaris10.props` file was modified with these parameters:

Note - The `root-password` string is copied from the source system's `/etc/shadow` file.

```
com.oracle.solaris.system.computer-name=TargetGuestDom

com.oracle.solaris.system.time-zone=US/Pacific

com.oracle.solaris.root-password=$5$jID2GKBG$RYoO3DUPeU4zbqwp3bOOwNsKoWlc/
i2GMKOTRHsSx90

com.oracle.solaris.system.ifname=net0

com.oracle.solaris.network.ipaddr.0=192.0.2.104

com.oracle.solaris.network.netmask.0=255.255.254.0

com.oracle.solaris.network.gateway.0=192.0.2.1

com.oracle.solaris.network.dns-servers.0="
198.51.100.1,198.51.100.2,198.51.100.3"

com.oracle.solaris.network.dns-search-domains.0="us.example.com,
examplecorp.com,example.com"

com.oracle.solaris.network.name-service=none
```

b. Configure and start the guest domain with the `ovmtconfig` utility.

This utility takes a couple of minutes to complete. For additional information refer to the *Oracle VM Server for SPARC Template Examples* at https://docs.oracle.com/cd/E62357_01/html/E62358/templateexamples.html.

Syntax:

`/opt/ovmtutils/bin/ovmtconfig -d` *domain-name* `-c` *commands* `-s -v -P` *Properties_file*

Where:

- `-d` *domain-name* – The name of the guest domain that will be configured by this command
- `-c` *commands* – For this example, specify the name of a script that contains commands that configure an Oracle Solaris 10 domain. The ovmtutils script called `ovmt_s10_sysidcfg.sh` and `ovmt_s11_expand_disk.sh` provide the correct commands for this option.
- `-s` – Starts the domain after configuration.
- `-v` – Create logical domain variables from properties.
- `-P` *Properties_file* – The path name of the properties file that contains configuration parameters (created in the previous step).

Example (entered on a single line):

```
root@TargetControlDom# /opt/ovmtutils/bin/ovmtconfig -s -d solaris10 -c /opt/
ovmtutils/share/scripts/ovmt_s11_expand_disk.sh,/opt/ovmtutils/
share/scripts/ovmt_s10_sysidcfg.sh -v -P /ovas1/solaris10.props
```

5. Verify the existence of the new target guest domain.

This example shows that the `solaris10` guest domain is created.

```
root@TargetControlDom# ldm ls
NAME             STATE      FLAGS     CONS       VCPU    MEMORY    UTIL   NORM   UPTIME
primary          active     -n-cv-    UART       16      32G       1.2%   1.2%   6d 14h
 11m
solaris10        active     -n----    5000       80      128G      0.1%   0.1%   0d 0h
 4m
```

6. Patch the target guest domain.

See "Patch the Target Guest Domain" on page 43.

Patch the Target Guest Domain

This procedure describes how to install required patches in the target guest domain. The patches were downloaded to the shared storage in "Prepare the Shared Storage" on page 35.

1. **Ensure that you are logged into the target guest domain as superuser.**

2. **Mount the shared storage.**

```
root@TargetGuestDom# mkdir /ovas1
root@TargetGuestDom# mount -F nfs TargetControlDom:/ovas1 /ovas1
```

3. **Ensure that the Flash archive patch 119534 is up to date on the target guest domain.**

   a. **Check if 119534 is up to date.**

   It must be 119534-33 (or later).

```
root@TargetGuestDom# showrev -p |grep 119534
Patch: 119534-33 Obsoletes: Requires: 119252-18, 120199-09, 126677-02 Incompatibles:
 Packages: SUNWinst
```

   b. **If the version of the patch is lower than 119534-33, install the patch.**

   In this example, the patch was downloaded and unzipped on the shared storage in "Prepare the Shared Storage" on page 35.

```
root@TargetGuestDom# cd /ovas1/Downloads
root@TargetGuestDom# patchadd ./119534-33
```

4. **Install the Oracle Solaris 10 recommended patch set on the target guest domain.**

   Install the patch set using the `installpatchset` command and include the required `s10patchset` passcode.

   In this example, the `installpatchset` command and the recommended patches are installed from the shared storage.

Note - At the end of the `installpatchset` script, this message is displayed: `Installation of patch set complete. PLEASE REBOOT THE SYSTEM`. Do not reboot the system at this time. Wait until instructed to reboot later in this procedure. Rebooting now will hinder your ability to perform some subsequent steps.

```
root@TargetGuestDom# cd /ovas1/Downloads/10_Recommended_CPU_2018-10
```

```
root@TargetGuestDom# ./installpatchset --s10patchset
```

5. Install the Oracle VM Server for SPARC 3.2 packages in the target guest domain.

The `20997186` package, once unzipped, provides the `SUNWldm` and `SUNWldmp2v` packages.

In this example, the `20997186` package was downloaded and unzipped on the shared storage in "Prepare the Shared Storage" on page 35.

a. Change to the `Product` directory that was created when the `20997186` package was unzipped.

Verify that the `SUNWldm.v` and `SUNWldmp2v` packages are present.

```
root@TargetGuestDom# cd /ovas1/Downloads/OVM_Server_SPARC-3_2/Product
root@TargetGuestDom# pwd
/ovas1/Downloads/OVM_Server_SPARC-3_2/Product
root@TargetGuestDom# ls
Japanese    SUNWldm.v   SUNWldmib   SUNWldmp2v
```

b. Use `pkgadd` to install the `SUNWldm` package.

Ignore the warnings and type `y` to continue the installation.

```
root@TargetGuestDom# pkgadd -Gd . SUNWldm.v
Processing package instance <SUNWldm.v> from </ovas1/Downloads/OVM_Server_SPARC-3_2/
Product>
LDoms Manager software(sparc.sun4v) 3.2.0.0.44,REV=2015.02.20.08.28
Copyright (c) 2015, Oracle and/or its affiliates. All rights reserved.
Using </> as the package base directory.
## Processing package information.
## Processing system information.
   15 package pathnames are already properly installed.
## Verifying package dependencies.
WARNING:
    The <SUNWldomr> package "Solaris Logical Domains
    (Root)" is a prerequisite package and should be
    installed.
WARNING:
    The <SUNWldomu> package "Solaris Logical Domains (Usr)"
    is a prerequisite package and should be installed.
Do you want to continue with the installation of <SUNWldm> [y,n,?] y
Installation of <SUNWldm> was successful.
```

c. Use `pkgadd` to install the `SUNWldmp2v` package.

```
root@TargetGuestDom# pkgadd -Gd . SUNWldmp2v
Processing package instance <SUNWldmp2v> from </ovas1/Downloads/
OVM_Server_SPARC-3_2/Product>
```

```
LDoms P2V tool(sparc) 3.2.0.0.44,REV=2015.02.20.08.28
Copyright (c) 2015, Oracle and/or its affiliates. All rights reserved.
Using </> as the package base directory.
## Processing package information.
## Processing system information.
   6 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

Installing LDoms P2V tool as <SUNWldmp2v>

## Installing part 1 of 1.
/opt/SUNWldmp2v/bin/ldmp2v_prepare
/opt/SUNWldmp2v/etc/ldmp2v.conf.sample
/opt/SUNWldmp2v/lib/ldmp2v_convert
/opt/SUNWldmp2v/lib/ldmp2v_finish
/opt/SUNWldmp2v/man/man1m/ldmp2v.1m <symbolic link>
/usr/sbin/ldmp2v
/usr/share/man/man1m/ldmp2v.1m
[ verifying class <none> ]
Installation of <SUNWldmp2v> was successful.
```

6. Install the 151934-06 (or later) patch on the target guest domain.

In this example, the patch is installed from the shared storage.

```
root@TargetGuestDom# cd /ovas1/Downloads
root@TargetGuestDom# patchadd ./151934-06
Patch packages installed:
  SUNWldm
  SUNWldmp2v
```

7. Verify the presence of the `ldmp2vz_convert` utility.

```
root@TargetGuestDom# ls /usr/sbin/ldmp2vz_convert
ldmp2vz_convert
```

8. Reboot the target guest domain.

```
root@TargetGuestDom# shutdown -i 6
```

9. Mount the shared storage.

The shared storage must be mounted after rebooting the system because access to the shared storage is needed in subsequent procedures.

```
root@TargetGuestDom# mount -F nfs TargetControlDom:/ovas1 /ovas1
```

Prepare the Source System

Perform this procedure on the Oracle Solaris 10 source system.

1. **Ensure that running transactions and connections are either left to complete or terminated based on business considerations.**

2. **Backup any critical data before the start of this process, so that you can fall back to this system if anything goes wrong.**

3. **Ensure that the Flash archive patch 119534 is up to date on the source system's global zone.**

   a. **Check if 119534 is up to date.**

   It must be 119534-33 (or later).

   ```
   root@SourceGlobal# showrev -p |grep 119534
   Patch: 119534-33 Obsoletes: Requires: 119252-18, 120199-09, 126677-02 Incompatibles:
    Packages: SUNWinst
   ```

   b. **If the version of the patch is lower than 119534-33, install the patch.**

   In this example, the patch was downloaded and unzipped on the shared storage in "Prepare the Shared Storage" on page 35.

   ```
   root@SourceGlobal# cd /ovas1/Downloads
   root@SourceGlobal# patchadd ./119534-33
   ```

4. **Install the Oracle VM Server for SPARC 3.2 packages in the source system's global zone.**

   The `20997186` package, once unzipped, provides the `SUNWldm` and `SUNWldmp2v` packages.

   In this example, the `20997186` package was downloaded and unzipped on the shared storage in "Prepare the Shared Storage" on page 35.

   a. **Change to the `Product` directory that was created when the `20997186` package was unzipped.**

   Verify that the `SUNWldm.v` and `SUNWldmp2v` packages are present.

   ```
   root@SourceGlobal# cd /ovas1/Downloads/OVM_Server_SPARC-3_2/Product
   root@SourceGlobal# pwd
   /ovas1/Downloads/OVM_Server_SPARC-3_2/Product
   root@SourceGlobal# ls
   Japanese    SUNWldm.v   SUNWldmib   SUNWldmp2v
   ```

b. Use `pkgadd` to install the `SUNWldm` package.

Ignore the warnings and type **`y`** to continue the installation.

```
root@SourceGlobal# pkgadd -Gd . SUNWldm.v
Processing package instance <SUNWldm.v> from </ovas1/Downloads/OVM_Server_SPARC-3_2/
Product>
LDoms Manager software(sparc.sun4v) 3.2.0.0.44,REV=2015.02.20.08.28
Copyright (c) 2015, Oracle and/or its affiliates. All rights reserved.
Using </> as the package base directory.
## Processing package information.
## Processing system information.
   15 package pathnames are already properly installed.
## Verifying package dependencies.
WARNING:
    The <SUNWldomr> package "Solaris Logical Domains
    (Root)" is a prerequisite package and should be
    installed.
WARNING:
    The <SUNWldomu> package "Solaris Logical Domains (Usr)"
    is a prerequisite package and should be installed.
Do you want to continue with the installation of <SUNWldm> [y,n,?] y
Installation of <SUNWldm> was successful.
```

c. Use `pkgadd` to install the `SUNWldmp2v` package.

```
root@SourceGlobal# pkgadd -Gd . SUNWldmp2v
Processing package instance <SUNWldmp2v> from </ovas1/Downloads/
OVM_Server_SPARC-3_2/Product>
LDoms P2V tool(sparc) 3.2.0.0.44,REV=2015.02.20.08.28
Copyright (c) 2015, Oracle and/or its affiliates. All rights reserved.
Using </> as the package base directory.
## Processing package information.
## Processing system information.
   6 package pathnames are already properly installed.
## Verifying package dependencies.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

Installing LDoms P2V tool as <SUNWldmp2v>

## Installing part 1 of 1.
/opt/SUNWldmp2v/bin/ldmp2v_prepare
/opt/SUNWldmp2v/etc/ldmp2v.conf.sample
/opt/SUNWldmp2v/lib/ldmp2v_convert
/opt/SUNWldmp2v/lib/ldmp2v_finish
```

```
/opt/SUNWldmp2v/man/man1m/ldmp2v.1m <symbolic link>
/usr/sbin/ldmp2v
/usr/share/man/man1m/ldmp2v.1m
[ verifying class <none> ]
Installation of <SUNWldmp2v> was successful.
```

5. **Install patch 151934-06 (or later) in the source system's global zone.**

   The 151934-06 patch provides the lift and shift utilities such as the `ldmp2vz_convert` utility.

   In this example, 151934-06 was downloaded and unzipped on the shared storage in "Prepare the Shared Storage" on page 35. It is installed from shared storage.

   a. **Use `patchadd` to install the patch.**

   ```
   root@SourceGlobal# cd /ovas1/Downloads
   root@SourceGlobal# patchadd 151934-06
   Patch packages installed:
     SUNWldm
     SUNWldmp2v
   ```

   b. **Verify the presence of the utilities.**

   ```
   root@SourceGlobal# ls /usr/sbin/ldmp2vz*
   ldmp2vz_collect      ldmp2vz_convert
   root@SourceGlobal # ls /opt/SUNWldm/lib/contrib/
   diskio      ovmtcreate  pigz        platinfo
   ```

Lifting and Shifting to the Target

These topics describe how to migrate the source system's environment, including the OS root file system and workloads, to an Oracle Solaris 10 guest domain on the target system:

▼ Lift the Source Environment to the Shared Storage

1. **Log into the source system and check the status of the workloads.**

   In this example, log into the `dbzone` as the `oracle1` user and check the status of the Database Listener.

```
-bash-3.2$ lsnrctl status

LSNRCTL for Solaris: Version 12.1.0.2.0 - Production on 30-Jul-2018 03:02:49

Copyright (c) 1991, 2014, Oracle.  All rights reserved.

Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=SourceDBzone)(PORT=1521)))
STATUS of the LISTENER
------------------------
Alias                     LISTENER
Version                   TNSLSNR for Solaris: Version 12.1.0.2.0 - Production
Start Date                29-JUL-2018 08:56:30
.........
.........
SNMP                      OFF
Listener Parameter File   /u01/app/oracle/product/12.1.0/grid/network/admin/listener.ora
Listener Log File         /u01/app/oracle/diag/tnslsnr/SourceDBzone/listener/alert/
log.xml
```

```
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=SourceDBzone)(PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=SourceDBzone)(PORT=5500))
(Security=(my_wallet_directory=/u01/app/oracle/product/12.1.0/dbhome_1/admin/orcl9/
xdb_wallet))(Presentation=HTTP)(Session=RAW))
Services Summary...
Service "+ASM" has 1 instance(s).
  Instance "+ASM", status READY, has 1 handler(s) for this service...
Service "mydb.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
Service "orcl9.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
Service "orcl9XDB.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
Service "pdborcl.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
The command completed successfully
```

2. **From a database client system, check the functionality of the database environment.**

   For example, issue a few SQL statements from the client to ensure connectivity and the contents.

```
root@client-sys# sqlplus system/welcome1@//SourceDBzone:1521/mydb.us.example.com
SQL> set pages 100
SQL> select * from v$recover_file;
no rows selected
SQL> select name from v$datafile union select member from v$logfile union select name
 from v$controlfile;
NAME
--------------------------------------------------------------------------------
+DATA/ORCL9/71335771C4785E92E054000B5DDC023F/DATAFILE/soe.278.981702033
+DATA/ORCL9/71335771C4785E92E054000B5DDC023F/DATAFILE/sysaux.274.981701921
+DATA/ORCL9/71335771C4785E92E054000B5DDC023F/DATAFILE/system.275.981701921
+DATA/ORCL9/71335771C4785E92E054000B5DDC023F/DATAFILE/users.277.981701955
+DATA/ORCL9/CONTROLFILE/current.261.981701477
+DATA/ORCL9/CONTROLFILE/current.262.981701475
+DATA/ORCL9/DATAFILE/undotbs1.260.981701345
/logs/redologs/redo04.log
/logs/redologs/redo05.log
/logs/redologs/redo06.log

10 rows selected.
```

3. **Sample one table in the application schema (SOE).**

   As a sanity check, the same table is sampled again after the migration to confirm data integrity.

```
SQL> select count(*) from soe.orders ;

  COUNT(*)
----------
   1429790
```

4. **Cleanly shutdown the Oracle Database, including any applications and third party products.**

   Shutdown the source `dbzone` database cleanly.

   ```
   -bash-3.2$ srvctl stop database -d orcl9
   ```

5. **Shutdown the Database.**

   This step prevents application related errors when the zone is started on the target.

   In this case, the application is the Oracle Database. The following commands also shutdown the database listener. After the deployment, `crsctl` is used to restart the high availability components.

   ```
   root@SourceDBzone# crsctl stop has
   root@SourceDBzone# crsctl disable has
   ```

6. **Verify the `webzone` Apache service.**

   a. **Log into `webzone` as superuser and check the state of the Apache web service.**

   ```
   root@SourceWebzone# svcs -a |grep apache2
   online         10:06:42 svc:/network/http:apache2
   root@SourceWebzone# ps -ef|grep apache
   webservd  8224  8181   0 03:19:24 ?           0:00 /usr/apache2/bin/httpd -k start
   webservd  8961  8181   0 03:47:50 ?           0:00 /usr/apache2/bin/httpd -k start
   ```

   b. **From an external client machine, run `wget index.html` to make sure the Apache Web Server is working.**

   ```
   root@client-sys# /usr/sfw/bin/wget http://SourceWebzone.com/
   --2018-07-17 03:53:23--  http://SourceWebzone.com/

   Resolving SourceWebzone.com (SourceWebzone.com)... 203.0.113.191
   Connecting to SourceWebzone.com (SourceWebzone.com)|203.0.113.191|:80... connected.
   HTTP request sent, awaiting response... 200 OK
   Length: 1456 (1.4K) [text/html]
   Saving to: 'index.html.2'
   ```

```
index.html.2        100%
[=============================================================================================
==================================>]   1.42K  --.-KB/s    in 0s

2018-07-17 03:53:23 (64.9 MB/s) - 'index.html.2' saved [1456/1456]
```

7. **From the global zone, shutdown and detach the non-global zones.**

   Failure to detach the non-global zones on the source system prevents the archive utility (`ldmp2vz_collect`) from completing.

```
root@SourceGlobal# zoneadm -z dbzone halt
root@SourceGlobal# zoneadm -z dbzone detach
root@SourceGlobal# zoneadm -z webzone halt
root@SourceGlobal# zoneadm -z webzone detach
```

8. **List the zones.**

```
root@SourceGlobal# zoneadm list -cv
  ID NAME             STATUS     PATH                           BRAND    IP
   0 global           running    /                              native   shared
   - webzone          configured /rpool/webzone                 native   shared
   - dbzone           configured /zones/dbzone                  native   shared
```

9. **Create a ZFS snapshot of the source system's zpools for `dbzone` and database binaries.**

   The `dbzone` zone path is installed on a zpool called `dbzone`. The Database and ASM binaries are installed on a zpool called `dbzone_db_binary`. The snapshots are later used to extract the `dbzone` and database binaries on the target system.

   Note - The `webzone` non-global zone has all its contents in the FLAR archive, therefore, it does not require any special techniques to transfer it from the source to the target.

```
root@SourceGlobal# zfs snapshot -r dbzone@first
root@SourceGlobal# zfs snapshot -r dbzone_db_binary@first

root@SourceGlobal# zfs list -t snapshot
NAME                     USED     AVAIL   REFER   MOUNTPOINT
dbzone@first             0        -       16.8G   -
dbzone_db_binary@first   0        -       21.5G   -
```

10. **Send the ZFS snapshots to the shared storage.**

    The `zfs send` command is used to transfer the snapshots from the source to shared storage. The `pigz` command reduces execution time by performing file compression using multiple threads in parallel.

```
root@SourceGlobal# zfs send dbzone@first | /opt/SUNWldm/lib/contrib/pigz > /ovas1/
dbzone.gz
root@SourceGlobal# zfs send dbzone_db_binary@first | /opt/SUNWldm/lib/contrib/pigz
 > /ovas1/dbzone_db_binary.gz
```

11. Make a binary copy of the application data.

Tip - Open two separate terminal windows on the source global zone and issue the following commands for each application disk.

---

a. Create a backup of the ASM data.

When transferring raw disk data, you can use the `dd` command to copy a slice, or `diskio` to copy the whole disk.

In this example, the ASM data occupies slice 0, so the `dd` command is used to copy the contents of that slice. The data can be compressed using `gzip`; in this case we use `pigz` which is a faster, parallel implementation of `gzip`.

```
root@SourceGlobal# dd if=/dev/rdsk/c2t600144F0E635D8C700005AC56B080015d0s0
 bs=104857600 | /opt/SUNWldm/lib/contrib/pigz > /ovas1/asm1.img.gz
root@SourceGlobal# dd if=/dev/rdsk/c2t600144F0E635D8C700005AC56B2E0016d0s0
 bs=104857600 | /opt/SUNWldm/lib/contrib/pigz > /ovas1/asm2.img.gz
```

b. Create a backup of the redo and archive logs.

```
root@SourceGlobal# ufsdump 0cf - /dev/md/rdsk/d20 | /opt/SUNWldm/lib/contrib/pigz
 > /ovas1/redo.ufsdump.gz
  DUMP: Date of this level 0 dump: Mon Jul 30 11:09:29 2018
  DUMP: Date of last level 0 dump: the epoch
  DUMP: Dumping /dev/md/rdsk/d20 (SourceGlobal:/zones/dbzone/root/logs/redologs) to
 standard output.
  DUMP: Mapping (Pass I) [regular files]
  DUMP: Mapping (Pass II) [directories]
  DUMP: Writing 63 Kilobyte records
  DUMP: Estimated 1232116 blocks (601.62MB).
  DUMP: Dumping (Pass III) [directories]
  DUMP: Dumping (Pass IV) [regular files]
  DUMP: 1232026 blocks (601.58MB) on 1 volume at 6315 KB/sec
  DUMP: DUMP IS DONE

root@SourceGlobal# ufsdump 0cf - /dev/md/rdsk/d30 | /opt/SUNWldm/lib/contrib/pigz
 > /ovas1/archive.ufsdump.gz
  DUMP: Date of this level 0 dump: Mon Jul 30 11:13:20 2018
  DUMP: Date of last level 0 dump: the epoch
```

```
 DUMP: Dumping /dev/md/rdsk/d30 (SourceGlobal:/zones/dbzone/root/logs/archivelogs)
to standard output.
 DUMP: Mapping (Pass I) [regular files]
 DUMP: Mapping (Pass II) [directories]
 DUMP: Writing 63 Kilobyte records
 DUMP: Estimated 1651844 blocks (806.56MB).
 DUMP: Dumping (Pass III) [directories]
 DUMP: Dumping (Pass IV) [regular files]
 DUMP: 1651858 blocks (806.57MB) on 1 volume at 7860 KB/sec
 DUMP: DUMP IS DONE
```

c. Verify that the backups are starting to show up in the shared storage.

```
root@SourceGlobal# cd /ovas1
root@SourceGlobal# ls -rtlh  *.gz
-rw-r--r--   1 root     root         11G Jul 30 07:58 dbzone.gz
-rw-r--r--   1 root     root        4.8G Jul 30 08:08 dbzone_db_binary.gz
-rw-r--r--   1 root     root        5.8G Jul 30 09:43 asm1.img.gz
-rw-r--r--   1 root     root        5.8G Jul 30 09:45 asm2.img.gz
-rw-r--r--   1 root     root         59M Jul 30 11:09 redo.ufsdump.gz
-rw-r--r--   1 root     root        185M Jul 30 11:14 archive.ufsdump.gz
```

12. Create a flash archive (FLAR) of the source file system.

The `ldmp2vz_collect` command creates a flash archive image (FLAR) of the source file system based on the configuration information that it collects about the source system.

Syntax:

`ldmp2vz_collect -c -d` *FLAR_directory*

Where:

- `-c` compresses the `.flar` file.
- `-d` *FLAR_directory* is the location where the `.flar` file is created.

In this example, the command is performed on the source system, but the `.flar` file is created on the shared storage making it available to the source and target systems.

Note - If you are not using shared storage, you can specify a local directory, then manually copy the `.flar` file to the target system.

```
root@SourceGlobal# ldmp2vz_collect -c -d /ovas1
INFO: Checking for patch 119534-33 (or higher)
INFO: Patch 119534-33 is installed, proceeding...

Backing up config files not retained by luupgrade process
```

```
Running: tar cvfX /etc/p2v/configs_tar /etc/p2v/configs_tar_exclude etc/defaultdomain
 etc/nsswitch.conf etc/pam.conf var/ldap
a etc/defaultdomain 1K
a etc/nsswitch.conf 1K
a etc/pam.conf 4K
a var/ldap/ 0K
a var/ldap/secmod.db 128K
a var/ldap/Oracle_SSL_CA_G2.pem 2K
a var/ldap/VTN-PublicPrimary-G5.pem 2K
a var/ldap/cachemgr.log 11K
a var/ldap/restore excluded
a var/ldap/key3.db 128K
a var/ldap/oracle_ssl_cert.pem 3K
a var/ldap/ldap_client_cred 1K
a var/ldap/ldap_client_file 3K
a var/ldap/cert8.db 64K

Running: flarcreate -n SourceGlobal -c -H -S -U
 content_architectures=sun4u,sun4v /ovas1/SourceGlobal.flar
Full Flash
Checking integrity...
Integrity OK.
Running precreation scripts...
Precreation scripts done.
Creating the archive...
Added keyword content_architectures=sun4u,sun4v does not begin with X-.
Archive creation complete.
Running postcreation scripts...
Postcreation scripts done.

Running pre-exit scripts...
Pre-exit scripts done.
flarcreate completed with exit code 0

Flar has been created in /ovas1/SourceGlobal.flar
-rw-r--r--   1 root     root     32350200956 Jul 30 13:14 /ovas1/SourceGlobal.flar
```

▼ Shift the Source Environment to the Target Guest Domain

This procedure deploys the FLAR on the target system.

1. **Log into the target guest domain.**

   The target guest domain name is `solaris10` in this example.

2. **Ensure that these items are available in the shared storage (`/ovas1`):**

   - FLAR (`SourceGlobal.flar`)
   - Oracle Solaris 10 ISO image (`sol-10-u11-ga-sparc-dvd.iso`)

```
root@SourceGlobal# cd /ovas1
root@SourceGlobal# ls -rtlh  *.gz  *.flar
-rw-r--r--   1 root     root         11G Jul 30 07:58 dbzone.gz
-rw-r--r--   1 root     root        4.8G Jul 30 08:08 dbzone_db_binary.gz
-rw-r--r--   1 root     root        5.8G Jul 30 09:43 asm1.img.gz
-rw-r--r--   1 root     root        5.8G Jul 30 09:45 asm2.img.gz
-rw-r--r--   1 root     root         59M Jul 30 11:09 redo.ufsdump.gz
-rw-r--r--   1 root     root        185M Jul 30 11:14 archive.ufsdump.gz
-rw-r--r--   1 root     root        32G  Jul 30 13:14 SourceGlobal.flar
```

3. **Ensure that the source database zone (`dbzone`) zpools and application data are copied over to the target system using `zfs send`.**

   Refer to instructions in "Lift the Source Environment to the Shared Storage" on page 49.

4. **Deploy the source image in the target Oracle Solaris 10 guest domain.**

   The `ldmp2vz_convert` command restores the source OS ZFS root file system to the target guest domain in a new boot environment. This restores the zone configuration and updates the new boot environment with SVR4 packages that are required for the sun4v target system.

   Syntax:

   `ldmp2vz_convert -f` *FLAR_filename* `-i` *iso_filename* `-b` *boot_environment*

   Where:

   - `-f` *FLAR_filename* – Is the file name of the FLAR created in "Lift the Source Environment to the Shared Storage" on page 49.
   - `-i` *iso_filename* – Is the file name of the Oracle Solaris ISO file.
   - `-b` *boot_environment* – Is the name of the guest domain boot environment.

   Example:

```
root@TargetGuestDom# cd /ovas1
root@TargetGuestDom# ldmp2vz_convert -f SourceGlobal.flar -i ./Downloads/sol-10-u11-ga-sparc-dvd.iso -b mybe
Renaming /var/yp to /var/yp.pre_p2v
Running: lucreate -s - -n mybe -p rpool -o /var/log/ldmp2vz/1669/lucreate.log
See /var/log/ldmp2vz/1669/lucreate.log for detailed output
Determining types of file systems supported
Validating file system requests
Preparing logical storage devices
Preparing physical storage devices
```

```
Configuring physical storage devices
Configuring logical storage devices
INFORMATION: No BEs are configured on this system.
The device </dev/dsk/c0d0s0> is not a root device for any boot environment; cannot get
 BE ID.
Running: zpool set autoexpand=on rpool
Renaming zfs dataset rpool/export to rpool/export_pre_p2v
Running: luupgrade -f -n mybe -s /tmp/p2v.sCa4qd/iso_unbundled -a /ovas1/
SourceGlobal.flar -o /var/log/ldmp2vz/1669/luupgrade1.log
See /var/log/ldmp2vz/1669/luupgrade1.log for detailed output
67352 blocks
cannot open 'dbzone_db_binary': dataset does not exist
cannot open 'dbzone_db_binary': dataset does not exist
/a
INFO: Removing sun4u packages:
SUNWcakr
SUNWcar
SUNWced
SUNWcpc
SUNWcpr
SUNWdscpr
SUNWdscpu
SUNWefc
SUNWifb
SUNWjfb
SUNWkvm
SUNWm64
SUNWnxge
SUNWpfb
SUNWpstl
SUNWsckmu
SUNWstc

Removal of <SUNWcakr> was successful.
Removal of <SUNWcar> was successful.
Removal of <SUNWced> was successful.
Removal of <SUNWcpc> was successful.
Removal of <SUNWcpr> was successful.
Removal of <SUNWdscpr> was successful.
Removal of <SUNWdscpu> was successful.
Removal of <SUNWefc> was successful.
/a/dev/fb*: No such file or directory
Removal of <SUNWifb> was successful.
/a/dev/fb*: No such file or directory
Removal of <SUNWjfb> was successful.
Removal of <SUNWkvm> was successful.
Removal of <SUNWm64> was successful.
Removal of <SUNWnxge> was successful.
```

```
Removal of <SUNWpfb> was successful.
Removal of <SUNWpstl> was successful.
Removal of <SUNWsckmu> was successful.
Removal of <SUNWstc> was successful.
Running: luupgrade -u -n mybe -s /tmp/p2v.sCa4qd/iso_unbundled -k /tmp/no-autoreg -
o /var/log/ldmp2vz/1669/luupgrade2.log
See /var/log/ldmp2vz/1669/luupgrade2.log for detailed output
67352 blocks
cannot open 'dbzone_db_binary': dataset does not exist
/a
Extracting configuration files
x etc/defaultdomain, 11 bytes, 1 tape blocks
x etc/nsswitch.conf, 434 bytes, 1 tape blocks
x etc/pam.conf, 3454 bytes, 7 tape blocks
x var/ldap, 0 bytes, 0 tape blocks
x var/ldap/secmod.db, 131072 bytes, 256 tape blocks
x var/ldap/Oracle_SSL_CA_G2.pem, 1830 bytes, 4 tape blocks
x var/ldap/VTN-PublicPrimary-G5.pem, 1732 bytes, 4 tape blocks
x var/ldap/cachemgr.log, 10274 bytes, 21 tape blocks
x var/ldap/key3.db, 131072 bytes, 256 tape blocks
x var/ldap/oracle_ssl_cert.pem, 2134 bytes, 5 tape blocks
x var/ldap/ldap_client_cred, 204 bytes, 1 tape blocks
x var/ldap/ldap_client_file, 2935 bytes, 6 tape blocks
x var/ldap/cert8.db, 65536 bytes, 128 tape blocks
Disabling Oracle Configuration Manager service in /a/var/svc/manifest/application/
management/ocm.xml
Running: luumount mybe
Running: luactivate mybe
A Live Upgrade Sync operation will be performed on startup of boot environment <mybe>.
**********************************************************************

The target boot environment has been activated. It will be used when you
reboot. NOTE: You MUST NOT USE the reboot, halt, or uadmin commands. You
MUST USE either the init or the shutdown command when you reboot. If you
do not use either init or shutdown, the system will not boot using the
target BE.

**********************************************************************

In case of a failure while booting to the target BE, the following process
needs to be followed to fallback to the currently working boot environment:

1. Enter the PROM monitor (ok prompt).

2. Boot the machine to Single User mode using a different boot device
(like the Solaris Install CD or Network). Examples:

     At the PROM monitor (ok prompt):
```

```
     For boot to Solaris CD:  boot cdrom -s
     For boot to network:     boot net -s

3. Mount the Current boot environment root slice to some directory (like
/mnt). You can use the following commands in sequence to mount the BE:

     zpool import rpool
     zfs inherit -r mountpoint rpool/ROOT/s10s_u11wos_24a
     zfs set mountpoint=<mountpointName> rpool/ROOT/s10s_u11wos_24a
     zfs mount rpool/ROOT/s10s_u11wos_24a

4. Run <luactivate> utility with out any arguments from the Parent boot
environment root slice, as shown below:

     <mountpointName>/sbin/luactivate

5. luactivate, activates the previous working boot environment and
indicates the result.
6. umount /mnt
7. zfs set mountpoint=/ rpool/ROOT/s10s_u11wos_24a
8. Exit Single User mode and reboot the machine.

**********************************************************************

Modifying boot archive service
Activation of boot environment <mybe> successful.
Run:
   init 6
which will cause a reboot.
After reboot run:
   svcadm disable ocm
   svcadm enable ldap/client
Then install any desired patch(es). Follow instructions to reboot again
if necessary. After rebooting, wait until the system boots to
multi-user-server milestone.
Check this using the command: svcs multi-user-server
and look for state 'online'.

Then copy and restore any required additional filesystems, for example non-ZFS
 filesystems, from system SourceGlobal.
```

5. **Reboot the target guest domain.**

Note - This reboots the target guest domain with the OS image from the source system. There might be issues with drivers and SMF services that are visible on the target guest domain console. Ignore the issues at this time. The driver issues are resolved after performing the "Install the Oracle Solaris 10 Recommended Patch Set" on page 60 procedure. The SMF services issues are resolved in "Complete the Target Guest Domain Configuration" on page 61.

```
root@TargetGuestDom# shutdown -i 6
```

6. **Once rebooted, log into the target domain and disable the Oracle Configuration Manager service and enable the LDAP service.**

```
root@TargetGuestDom# svcadm disable ocm
root@TargetGuestDom# svcadm enable ldap/client
```

7. **Verify that the zones are present in the boot environment.**

```
root@TargetGuestDom# zoneadm list -cv
  ID NAME             STATUS     PATH                           BRAND    IP
   0 global           running    /                              native   shared
   - webzone          configured /rpool/webzone                 native   shared
   - dbzone           configured /zones/dbzone                  native   shared
```

8. **From the target guest domain, mount the shared storage.**

```
root@TargetGuestDom# mount -F nfs TargetControlDom:/ovas1 /ovas1
```

▼ Install the Oracle Solaris 10 Recommended Patch Set

Although the patch set was installed on the target guest domain in "Patch the Target Guest Domain" on page 43, the patch set must be installed again so that the newly created boot environment is properly patched. This ensures that the non-global zones are up to date with the latest required patches.

1. **Log into the target guest domain as superuser.**

2. **Install the `10_Recommended_CPU_2018-10` patch set on the target guest domain.**

   In this example, the `installpatchset` command and the recommended patches are installed from the shared storage.

```
root@TargetGuestDom# cd /ovas1/Downloads/10_Recommended_CPU_2018-10
root@TargetGuestDom# ./installpatchset --s10patchset
```

3. **Reboot the target guest domain.**

```
root@TargetGuestDom# shutdown -i 6
```

Complete the Target Guest Domain Configuration

This procedure describes how to make a few administrative changes to ensure that the target guest domain is properly configured in the target environment.

1. **Log into the target guest domain as superuser.**

2. **Mount the shared storage.**

```
root@TargetGuestDom# mount -F nfs TargetControlDom:/ovas1 /ovas1
```

3. **Adjust the swap and dump sizes on the target guest domain.**

It is possible that swap and dump sizes on the source system are larger than the migrated size of swap and dump on the target system. Refer to information collected in "Review the Source System Configuration" on page 21.

In this example, the swap and dump sized are adjusted to match the sizes from the source system to ensure these resources can properly support the workloads.

```
root@TargetGuestDom# zfs set volsize=160g rpool/swap
root@TargetGuestDom# zfs set volsize=3g rpool/dump
```

4. **Configure network parameters to suit the target network environment.**

The `ldmp2vz_convert` script copies these configuration files from the source system into the target guest domain:

- `/etc/defaultdomain`
- `/etc/nsswitch.conf`
- `/etc/pam.conf`
- `/var/ldap/`

If the network environment of the target guest domain matches that of the source system, the settings in the configuration files are likely suitable and do not need to be changed. In this case, skip the following substeps.

If the target guest domain environment differs from that of the source system, the configuration files might not function correctly on the target domain. For example, you might have executed the `ovmtconfig` command to set up the target guest domain. If this is the case, you can restore the configuration files to their default contents in the target guest domain by temporarily

mounting the original boot environment and using the `tar` command to copy the files from there to the current boot environment.

a. List the status of the target boot environments.

In this example, `s10s_u11wos_24a` is the original target boot environment that has the proper configuration files. `mybe` is the boot environment created by the `ldmp2vz_convert` command (in "Shift the Source Environment to the Target Guest Domain" on page 55) and has configuration files from the source system.

```
root@TargetGuestDom# lustatus
Boot Environment           Is       Active Active    Can    Copy
Name                       Complete Now    On Reboot Delete Status
-------------------------- -------- ------ --------- ------ ----------
s10s_u11wos_24a            yes      no     no        yes    -
mybe                       yes      yes    yes       no     -
```

b. Create a directory and move the configuration files to it.

This step moves the unwanted files to a safe location in the unlikely event that they need to be referenced for some reason.

```
root@TargetGuestDom# cd /
root@TargetGuestDom# mkdir /etc/p2v/configs_backup
root@TargetGuestDom# mv etc/defaultdomain etc/nsswitch.conf etc/pam.conf var/
ldap /etc/p2v/configs_backup
```

c. Mount the original boot environment to a temporary mount point on the guest domain.

```
root@TargetGuestDom# lumount s10s_u11wos_24a /mnt
/mnt
```

d. In the original boot environment, create a tar file of the configuration files.

```
root@TargetGuestDom# cd /mnt
root@TargetGuestDom# tar cvf /etc/p2v/configs_tar_orig etc/defaultdomain etc/
nsswitch.conf etc/pam.conf var/ldap
a etc/defaultdomain 1K
a etc/nsswitch.conf 2K
a etc/pam.conf 4K
a var/ldap/ 0K
a var/ldap/ldap_client_file 3K
a var/ldap/ldap_client_cred 1K
a var/ldap/cachemgr.log 4K
a var/ldap/restore/ 0k
```

e. Extract the tar file in the target guest domain.

```
root@TargetGuestDom# cd /
root@TargetGuestDom# tar xvf /etc/p2v/configs_tar_orig
x etc/defaultdomain, 11 bytes, 1 tape blocks
x etc/nsswitch.conf, 1202 bytes, 3 tape blocks
x etc/pam.conf, 3263 bytes, 7 tape blocks
x var/ldap, 0 bytes, 0 tape blocks
x var/ldap/ldap_client_file, 2935 bytes, 6 tape blocks
x var/ldap/ldap_client_cred, 204 bytes, 1 tape blocks
x var/ldap/cachemgr.log, 3469 bytes, 7 tape blocks
x var/ldap/restore, 0 bytes, 0 tape blocks
```

5. Consider reconfiguring other services that have a different configuration in the target environment.

There might be other services whose configuration has to be changed in the target guest domain rather than using the source machine's configuration. For example, after rebooting the target guest domain, you might see `syslog` messages like this example:

```
Jul  30 14:40:40 TargetGuestDom ntpdate[415]: no server suitable for synchronization
 found
```

This message occurs because the NTP client configuration file was restored by `ldmp2vz_convert` from the source system. Refer to the xntpd(1M) for further information on NTP configuration.

Some SMF services carried over from the source system OS might not run because the platform is different, for example during a sun4u to sun4v migration. In this case, disable the SMF services, for example:

```
root@TargetGuestDom# svcadm disable dcs
root@TargetGuestDom# svcadm disable oplhpd
root@TargetGuestDom# svcadm disable sckmd
```

▼ Configure Target Domain Virtual Disks

In this procedure, several vdisks are added to the target guest domain to support the incoming zones and the application data.

As you perform this procedure, take into account the state of your target system and adjust or omit steps as needed.

1. Add all the disks to the target guest domain.

This is an example of adding one vdisk, but the commands must be repeated to add all the disks that were provisioned in "Prepare the Target System" on page 29.

a. Log into the target control domain as superuser.

b. Add a disk to the target guest domain.

```
root@TargetControlDom# ldm add-vdsdev /dev/rdsk/
c0t600144F09F2C0BFD00005BE4A90F0004d0s2 solaris10-vol6@ovmt-vds0
root@TargetControlDom# ldm add-vdisk vdisk6 solaris10-vol6@ovmt-vds0 solaris10
```

c. Repeat Step 1b to add each disk to the target guest domain.

d. Verify the disk configuration.

```
root@TargetControlDom# ldm list -o disk primary

NAME
primary

VDS
    NAME         VOLUME         OPTIONS          MPGROUP        DEVICE
    ovmt-vds0    solaris10-vol0                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4A8500003d0s2
                 solaris10-vol1                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4A90F0004d0s2
                 solaris10-vol2                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4A9A90005d0s2
                 solaris10-vol3                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4BF670006d0s2
                 solaris10-vol4                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4C4BB000Dd0s2
                 solaris10-vol5                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4C4E4000Ed0s2
                 solaris10-vol6                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4BFC90007d0s2
                 solaris10-vol7                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4BFE60008d0s2
                 solaris10-vol8                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4C0410009d0s2
                 solaris10-vol9                                 /dev/rdsk/
c0t600144F09F2C0BFD00005BE4C06B000Ad0s2
                 solaris10-vol10                                /dev/rdsk/
c0t600144F09F2C0BFD00005BE4C3EC000Bd0s2
                 solaris10-vol11                                /dev/rdsk/
c0t600144F09F2C0BFD00005BE4C424000Cd0s2
```

2. List the vdisk bindings.

The vdisks assigned to the target guest domain are `vdisk0` through `vdisk10`, however when you display the disks with the `format` command, they are listed as `c0d0` through `c0d10`.

```
root@TargetControlDom# ldm ls -o disk solaris10
NAME
solaris10

DISK
    NAME         VOLUME                 TOUT ID   DEVICE  SERVER         MPGROUP
    vdisk0       solaris10-vol0@ovmt-vds0      0    disk@0  primary

    vdisk1       solaris10-vol1@ovmt-vds0      1    disk@1  primary

    vdisk2       solaris10-vol2@ovmt-vds0      2    disk@2  primary

    vdisk3       solaris10-vol3@ovmt-vds0      3    disk@3  primary

    vdisk4       solaris10-vol4@ovmt-vds0      4    disk@4  primary

    vdisk5       solaris10-vol5@ovmt-vds0      5    disk@5  primary

    vdisk6       solaris10-vol6@ovmt-vds0      6    disk@6  primary

    vdisk7       solaris10-vol7@ovmt-vds0      7    disk@7  primary

    vdisk8       solaris10-vol8@ovmt-vds0      8    disk@8  primary

    vdisk9       solaris10-vol9@ovmt-vds0      9    disk@9  primary

    vdisk10      solaris10-vol10@ovmt-vds0      10   disk@10 primary

    vdisk11      solaris10-vol11@ovmt-vds0      11   disk@11 primary
```

3. Log into the target guest domain and verify the target guest domain disks.

```
root@TargetGuestDom# echo|format

Searching for disks...done
c0d1: configured with capacity of 599.92GB
c0d2: configured with capacity of 299.91GB
c0d3: configured with capacity of 299.91GB
c0d4: configured with capacity of 149.91GB
c0d5: configured with capacity of 149.91GB
c0d6: configured with capacity of 199.93GB
c0d7: configured with capacity of 199.93GB
c0d8: configured with capacity of 249.92GB
c0d9: configured with capacity of 249.92GB
```

```
c0d10: configured with capacity of 1023.75MB
c0d11: configured with capacity of 1023.75MB


AVAILABLE DISK SELECTIONS:
       0. c0d0 <SUN-DiskImage-16GB cyl 17064 alt 2 hd 96 sec 768>
          /virtual-devices@100/channel-devices@200/disk@0
       1. c0d1 <SUN-ZFSStorage7355-1.0 cyl 19501 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@1
       2. c0d2 <SUN-ZFSStorage7355-1.0 cyl 9749 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@2
       3. c0d3 <SUN-ZFSStorage7355-1.0 cyl 9749 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@3
       4. c0d4 <SUN-ZFSStorage7355-1.0 cyl 4873 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@4
       5. c0d5 <SUN-ZFSStorage7355-1.0 cyl 4873 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@5
       6. c0d6 <SUN-ZFSStorage7355-1.0 cyl 6499 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@6
       7. c0d7 <SUN-ZFSStorage7355-1.0 cyl 6499 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@7
       8. c0d8 <SUN-ZFSStorage7355-1.0 cyl 8124 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@8
       9. c0d9 <SUN-ZFSStorage7355-1.0 cyl 8124 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@9
      10. c0d10 <SUN-ZFSStorage7355-1.0 cyl 8190 alt 2 hd 8 sec 32>
          /virtual-devices@100/channel-devices@200/disk@a
      11. c0d11 <SUN-ZFSStorage7355-1.0 cyl 8190 alt 2 hd 8 sec 32>
          /virtual-devices@100/channel-devices@200/disk@b
Specify disk (enter its number): Specify disk (enter its number):
```

4. Prepare the disk for the mirror rpool.

a. Identify the rpool disk.

In this example, the rpool disk is `c0d0s0`.

```
root@TargetGuestDom# zpool status rpool
  pool: rpool
 state: ONLINE
config:

        NAME STATE READ WRITE CKSUM
        rpool ONLINE 0 0 0
          c0d0s0 ONLINE 0 0 0

errors: No known data errors
```

b. On the target guest domain, run the `prtvtoc` command on the rpool disk to collect the following disk geometry information:

Values for this example are shown in parenthesis.

- Sectors per track (768)
- Tracks per cylinder (96)
- Accessible cylinders (17064)

These values are used to configure the disk geometry of the rpool mirror disk so that the two disks match.

```
root@TargetGuestDom# prtvtoc /dev/rdsk/c0d0s2
* /dev/rdsk/c0d1s2 partition map
*
* Dimensions:
*     512 bytes/sector
*     768 sectors/track
*      96 tracks/cylinder
*   73728 sectors/cylinder
*   17066 cylinders
*   17064 accessible cylinders
*
* Flags:
*   1: unmountable
*  10: read-only
*
*                          First     Sector    Last
* Partition  Tag  Flags    Sector     Count    Sector  Mount Directory
       0      2    00          0 1258094592 1258094591
       2      5    00          0 1258094592 1258094591
```

c. Use the `format` utility to configure the mirror disk with the same disk geometry information.

```
root@TargetGuestDom# format c0d1

selecting c0d1

[disk formatted]
FORMAT MENU:
        disk       - select a disk
        type       - select (define) a disk type
        partition  - select (define) a partition table
        current    - describe the current disk
        format     - format and analyze the disk
```

```
        repair     - repair a defective sector
        label      - write label to the disk
        analyze    - surface analysis
        defect     - defect list management
        backup     - search for backup labels
        verify     - read and display labels
        save       - save new disk/partition definitions
        inquiry    - show disk ID
        volname    - set 8-character volume name
        !<cmd>     - execute <cmd>, then return
        quit

format> type

AVAILABLE DRIVE TYPES:
        0. Auto configure
        1. Quantum ProDrive 80S
        2. Quantum ProDrive 105S
        3. CDC Wren IV 94171-344
        4. SUN0104
        5. SUN0207
        6. SUN0327
        7. SUN0340
        8. SUN0424
        9. SUN0535
        10. SUN0669
        11. SUN1.0G
        12. SUN1.05
        13. SUN1.3G
        14. SUN2.1G
        15. SUN2.9G
        16. Zip 100
        17. Zip 250
        18. Peerless 10GB
        19. SUN-ZFS Storage 7355-1.0
        20. other
Specify disk type (enter its number)[19]: 20
Enter number of data cylinders: 17064
Enter number of alternate cylinders[2]:
Enter number of physical cylinders[17066]:
Enter number of heads: 96
Enter physical number of heads[default]:
Enter number of data sectors/track: 768
Enter number of physical sectors/track[default]:
Enter rpm of drive[3600]:
Enter format time[default]:
Enter cylinder skew[default]:
Enter track skew[default]:
```

```
Enter tracks per zone[default]:
Enter alternate tracks[default]:
Enter alternate sectors[default]:
Enter cache control[default]:
Enter prefetch threshold[default]:
Enter minimum prefetch[default]:
Enter maximum prefetch[default]:
Enter disk type name (remember quotes): "600GB disk for guest second root disk"
selecting c0d1
[disk formatted]
format> l
Ready to label disk, continue? y
format> q
```

d. Verify that the disk geometry for the rpool and mirror match.

```
root@TargetGuestDom# prtvtoc /dev/rdsk/c0d1s2
* /dev/rdsk/c0d1s2 partition map
*
* Dimensions:
*     512 bytes/sector
*     768 sectors/track
*      96 tracks/cylinder
*   73728 sectors/cylinder
*   17066 cylinders
*   17064 accessible cylinders
*
* Flags:
*   1: unmountable
*  10: read-only
*
*                          First     Sector    Last
* Partition  Tag  Flags    Sector     Count    Sector  Mount Directory
       0      2    00          0    294912    294911
       1      3    01     294912    294912    589823
       2      5    01          0 1258094592 1258094591
       6      4    00     589824 1257504768 1258094591
```

e. Update the mirror disk VTOC.

```
root@TargetGuestDom# prtvtoc /dev/rdsk/c0d0s2 | fmthard -s - /dev/rdsk/c0d1s2
fmthard:  New volume table of contents now in place.
```

5. Prepare the LUNs used by ASM.

These substeps describe how to configure the ASM target disks with the same disk label type and partition table as the ASM disks on the source system.

a. On the source system, use `prtvtoc` to capture the partition table of the LUNs used by ASM.

```
root@SourceGlobal# prtvtoc /dev/rdsk/c2t600144F0E635D8C700005AC56B080015d0s2
 > /ovas1/dbzone_asm_disk1.txt
```

b. Review that partition table to determine if it has an EFI or SMI label.

In the `Dimensions` list, SMI labels describe the disk geometry with cylinders. EFI labels do not specify cylinders, as shown in the following example.

```
root@TargetGuestDom# cat /ovas1/dbzone_asm_disk1.txt
* /dev/rdsk/c2t600144F0E635D8C700005AC56B080015d0s2 partition map
*
* Dimensions:
*     512 bytes/sector
* 419430400 sectors
* 419430333 accessible sectors
*
* Flags:
*   1: unmountable
*  10: read-only
*
*                          First     Sector    Last
* Partition  Tag  Flags    Sector     Count    Sector  Mount Directory
       0      4    00          34 419413949 419413982
       8     11    00  419413983     16384 419430366
```

c. From the target guest domain, determine which disks will be used for ASM.

```
root@TargetGuestDom# ldm ls -o disk | egrep "c0t600144F09F2C0BFD00005BE4BFC90007d0|
c0t600144F09F2C0BFD00005BE4BFE60008d0"
                solaris10-vol6                                  /dev/rdsk/
c0t600144F09F2C0BFD00005BE4BFC90007d0s2
                solaris10-vol7                                  /dev/rdsk/
c0t600144F09F2C0BFD00005BE4BFE60008d0s2

root@TargetGuestDom# ldm ls -o disk solaris10 | egrep "solaris10-vol6|solaris10-
vol7"
    vdisk6       solaris10-vol6@ovmt-vds0      6    disk@6  primary

    vdisk7       solaris10-vol7@ovmt-vds0      7    disk@7  primary
```

d. For each ASM disk, use `format -e` to label the disk with the same label type as the source LUN.

```
root@TargetGuestDom# format -e c0d6
selecting c0d6
...
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 1

root@TargetGuestDom# format -e c0d7
selecting c0d6
...
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 1
```

e. Transfer the source ASM disk partition table to each target ASM disk.

Because the target disks for ASM are provisioned with the same size as the source disks, the source disk partition table can be transferred to the target disk partition table using the `fmthard` command.

```
root@TargetGuestDom# fmthard -s /ovas1/dbzone_asm_disk1.txt /dev/rdsk/c0d6s2
root@TargetGuestDom# fmthard -s /ovas1/dbzone_asm_disk1.txt /dev/rdsk/c0d7s2
```

6. Create a mirror for the target guest domain rpool.

```
root@TargetGuestDom# zpool attach rpool c0d0s0 c0d1s0
Make sure to wait until resilver is done before rebooting.
```

7. Configure metadevices and create UFS file systems.

The source `dbzone` uses two separate UFS file systems on SVM metadevices for the redo and archive logs, so the same metadevices are created on the target guest domain. One metadevice, `d20`, is created with 10 GB of storage for the Redo log file system. Another metadevice, `d30`, is created with 200 GB of storage for the Archive Redo log file system.

```
root@TargetGuestDom# metadb -a -c 3 -f /dev/dsk/c0d10s4
root@TargetGuestDom# metadb -a -c 3 -f /dev/dsk/c0d11s4

root@TargetGuestDom# metadb
        flags           first blk       block count
     a        u         16              8192            /dev/dsk/c0d10s4
     a        u         8208            8192            /dev/dsk/c0d10s4
     a        u         16400           8192            /dev/dsk/c0d10s4
     a        u         16              8192            /dev/dsk/c0d11s4
     a        u         8208            8192            /dev/dsk/c0d11s4
     a        u         16400           8192            /dev/dsk/c0d11s4
```

```
root@TargetGuestDom# metainit d11 1 1 c0d8s0
d11: Concat/Stripe is setup
root@TargetGuestDom# metainit d12 1 1 c0d9s0
d12: Concat/Stripe is setup
root@TargetGuestDom# metainit d0 -m d11 d12
metainit: d0: WARNING: This form of metainit is not recommended.
The submirrors may not have the same data.
Please see ERRORS in metainit(1M) for additional information.
d0: Mirror is setup

root@TargetGuestDom# metainit d20 -p d0 10g
d20: Soft Partition is setup
root@TargetGuestDom# metainit d30 -p d0 200g
d30: Soft Partition is setup

root@TargetGuestDom# metastat
d30: Soft Partition
    Device: d0
    State: Okay
    Size: 419430400 blocks (200 GB)
        Extent Start Block Block count
             0 21036096 419430400

d0: Mirror
    Submirror 0: d11
      State: Okay
    Submirror 1: d12
      State: Okay
    Pass: 1
    Read option: roundrobin (default)
    Write option: parallel (default)
    Size: 524127984 blocks (249 GB)

d11: Submirror of d0
    State: Okay
    Size: 524127984 blocks (249 GB)
    Stripe 0:
        Device Start Block Dbase State Reloc Hot Spare
        c0d8s0 0 No Okay Yes


d12: Submirror of d0
    State: Okay
    Size: 524127984 blocks (249 GB)
    Stripe 0:
        Device Start Block Dbase State Reloc Hot Spare
        c0d9s0 0 No Okay Yes
```

```
d20: Soft Partition
    Device: d0
    State: Okay
    Size: 20971520 blocks (10 GB)
        Extent Start Block Block count
             0 64544 20971520

Device Relocation Information:
Device Reloc Device ID
c0d9 Yes id1,vdc@n600144f09f2c0bfd00005be4c06b000a
c0d8 Yes id1,vdc@n600144f09f2c0bfd00005be4c0410009


root@TargetGuestDom# newfs /dev/md/rdsk/d20
newfs: construct a new file system /dev/md/rdsk/d20: (y/n)? y
Warning: 4096 sector(s) in last cylinder unallocated
/dev/md/rdsk/d20:       20971520 sectors in 3414 cylinders of 48 tracks, 128 sectors
        10240.0MB in 214 cyl groups (16 c/g, 48.00MB/g, 5824 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
 32, 98464, 196896, 295328, 393760, 492192, 590624, 689056, 787488, 885920,
 20055584, 20154016, 20252448, 20350880, 20449312, 20547744, 20646176,
 20744608, 20843040, 20941472

root@TargetGuestDom# newfs /dev/md/rdsk/d30
newfs: construct a new file system /dev/md/rdsk/d30: (y/n)? y
Warning: 2048 sector(s) in last cylinder unallocated /dev/md/rdsk/d30:       419430400
 sectors in 68267 cylinders of 48 tracks, 128 sectors
        204800.0MB in 4267 cyl groups (16 c/g, 48.00MB/g, 5824 i/g)
super-block backups (for fsck -F ufs -o b=#) at: 32, 98464, 196896, 295328, 393760,
 492192, 590624, 689056, 787488, 885920,
Initializing cylinder groups:
................................................................................
super-block backups for last 10 cylinder groups at:
 418484384, 418582816, 418681248, 418779680, 418878112, 418976544, 419074976,
 419173408, 419271840, 419370272
```

8. Create zpools for the `dbzone`.

```
root@TargetGuestDom# mkdir -p /zones/dbzone
root@TargetGuestDom# zpool create -m /zones/dbzone dbzone mirror c0d2 c0d3
root@TargetGuestDom# zpool create -f dbzone_db_binary mirror c0d4 c0d5
```

9. Verify the zpools.

```
root@TargetGuestDom# zpool status
  pool: dbzone
```

```
 state: ONLINE
 scan: none requested
config:
        NAME        STATE     READ WRITE CKSUM
        dbzone      ONLINE       0     0     0
          mirror-0  ONLINE       0     0     0
            c0d2    ONLINE       0     0     0
            c0d3    ONLINE       0     0     0
errors: No known data errors
  pool: dbzone_db_binary
 state: ONLINE
 scan: none requested
config:
        NAME        STATE     READ WRITE CKSUM
        dbzone_db_binary  ONLINE       0     0     0
          mirror-0  ONLINE       0     0     0
            c0d4    ONLINE       0     0     0
            c0d5    ONLINE       0     0     0
errors: No known data errors
  pool: rpool
 state: ONLINE
 scan: resilvered 11.5G in 0h1m with 0 errors on Mon Jul 30 15:02:20 2018
config:
        NAME        STATE     READ WRITE CKSUM
        rpool       ONLINE       0     0     0
          mirror-0  ONLINE       0     0     0
            c0d0s0  ONLINE       0     0     0
            c0d1s0  ONLINE       0     0     0
errors: No known data errors
```

Reconfiguring the Zones on the Target System

These topics describe how to complete the lift and shift by reconfiguring the zones and reviewing the migrated configuration in the new environment.

▼ Reconfigure the Zones

The target zones still contain configuration information from the source system (for example, network parameters) and must be adjusted for the new environment.

1. **Reconfigure the `dbzone` zone.**

```
root@TargetGuestDom# zonecfg -z dbzone
zonecfg:dbzone> remove net
zonecfg:dbzone> remove device
zonecfg:dbzone> commit
zonecfg:dbzone> add net
zonecfg:dbzone:net> set physical=vnet0
zonecfg:dbzone:net> set address=192.0.2.52/23
zonecfg:dbzone:net> end
zonecfg:dbzone> add device
zonecfg:dbzone:device> set match=/dev/rdsk/c0d2s0
zonecfg:dbzone:device> end
zonecfg:dbzone> add device
zonecfg:dbzone:device> set match=/dev/rdsk/c0d3s0
zonecfg:dbzone:device> end
zonecfg:dbzone> commit
zonecfg:dbzone> verify
```

```
zonecfg:dbzone> exit
```

2. Verify the `dbzone` configuration.

```
root@TargetGuestDom# zonecfg -z dbzone info
zonename: dbzone
zonepath: /zones/dbzone
brand: native
autoboot: true
bootargs:
pool: dbzone-pool
limitpriv: default,proc_priocntl,proc_clock_highres
scheduling-class:
ip-type: shared
hostid:
fs:
        dir: /logs/redologs
        special: /dev/md/dsk/d20
        raw: /dev/md/rdsk/d20
        type: ufs
        options:
 [rw,setuid,devices,intr,forcedirectio,largefiles,logging,noquota,xattr,nodfratime]
fs:
        dir: /logs/archivelogs
        special: /dev/md/dsk/d30
        raw: /dev/md/rdsk/d30
        type: ufs
        options:
 [rw,setuid,devices,intr,forcedirectio,largefiles,logging,noquota,xattr,nodfratime]
net:
        address: 192.0.2.52/23
        physical: vnet0
        defrouter not specified
device
        match: /dev/rdsk/c0d2s0
device
        match: /dev/rdsk/c0d3s0
dataset:
        name: dbzone_db_binary
```

3. Reconfigure the `webzone` zone.

```
root@TargetGuestDom# zonecfg -z webzone
zonecfg:webzone> remove net
zonecfg:webzone> add net
zonecfg:webzone:net> set address=192.0.2.26/23
zonecfg:webzone:net> set physical=vnet0
zonecfg:webzone:net> end
```

```
zonecfg:webzone> commit
zonecfg:webzone> verify
zonecfg:webzone> exit
```

4. Verify the `webzone` configuration.

```
root@TargetGuestDom# zonecfg -z webzone info
zonename: webzone
zonepath: /rpool/webzone
brand: native
autoboot: true
bootargs:
pool: webzone-pool
limitpriv:
scheduling-class:
ip-type: shared
hostid:
inherit-pkg-dir:
        dir: /usr
inherit-pkg-dir:
        dir: /sbin
net:
        address: 192.0.2.26/23
        physical: vnet0
        defrouter not specified
```

5. Transfer the data from the shared storage to the zpools.

```
root@TargetGuestDom# /opt/SUNWldm/lib/contrib/pigz -dc -f /ovas1/dbzone.gz | zfs receive
 -F dbzone@first
root@TargetGuestDom# /opt/SUNWldm/lib/contrib/pigz -dc -f /ovas1/dbzone_db_binary.gz |
 zfs receive -F dbzone_db_binary@first
```

6. Attach the zones.

Immediately after the zones are attached, the zones are in an installed state.

```
root@TargetGuestDom# zoneadm -z dbzone attach -u
Getting the list of files to remove
Removing 485 files
Remove 17 of 17 packages
Installing 8273 files
Add 87 of 87 packages
Updating editable files
The file </var/sadm/system/logs/update_log> within the zone contains a log of the zone
 update.

root@TargetGuestDom# zoneadm -z webzone attach -u
Getting the list of files to remove
```

```
Removing 279 files
Remove 17 of 17 packages
Installing 861 files
Add 87 of 87 packages
Updating editable files
The file </var/sadm/system/logs/update_log> within the zone contains a log of the zone
 update.
```

7. **Verify that the zones are in an installed state.**

```
root@TargetGuestDom# zoneadm list -cv

  ID NAME             STATUS         PATH                         BRAND    IP
   0 global           running          /                          native   shared
   - webzone          installed        /rpool/webzone             native   shared
   - dbzone           installed        /zones/dbzone              native   shared
```

Reconfigure Zone CPU IDs

At this point, the processor sets in the new target domain are configured for CPU IDs that correspond to the source system. This configuration does not correspond to the CPU IDs available in the target system. Use this procedure to reset the pool configuration and assign the appropriate CPU IDs to the recreated processor sets and pools. The order of the core allocation is configured so that it is similar to the core allocation of the source system.

For additional information about resource management on Oracle Solaris zones, refer to the administration guide at: https://docs.oracle.com/cd/E26505_01/html/817-1592/toc.html.

1. **Display the CPUs and CPU IDs that are available in the target domain.**
   In this example, the `psrinfo` command shows that the target guest domain has the following CPUs:

```
root@TargetGuestDom# psrinfo -pv
The physical processor has 16 virtual processors (0-15)
  SPARC-S7 (chipid 0, clock 4267 MHz)
The physical processor has 32 virtual processors (16-47)
  SPARC-S7 (chipid 1, clock 4267 MHz)
The physical processor has 32 virtual processors (48-79)
  SPARC-S7 (chipid 2, clock 4267 MHz)
```

2. **Reassign processors.**
   On the source, the `dbzone` was assigned the last 32 virtual processors, and the `webzone` was assigned the previous 32 virtual processors. Therefore, on the target system, a similar

configuration is applied where the dbzone is assigned the last 32 virtual processors (48-79) and the webzone is assigned processors 16-47.

The following commands are used to reassign the dbzone processors. The startcpu variable specifies the starting CPU ID , and the endcpu variable specifies the ending CPU ID.

```
TargetGuestDom# pooladm -x
TargetGuestDom# pooladm -c discover
TargetGuestDom# poolcfg -c 'create pset dbzone-set(uint pset.min=32;uint pset.max=32)'
TargetGuestDom# poolcfg -c "create pool dbzone-pool"
TargetGuestDom# poolcfg -c 'associate pool dbzone-pool (pset dbzone-set)'
TargetGuestDom# startcpu=48; endcpu=79
TargetGuestDom# i=$startcpu; while [ ${i} -le $endcpu ]; do set -x; poolcfg -c "transfer
 to pset dbzone-set ( cpu ${i} )"; set +x; i= $((${i} + 1)); done
```

The following commands are used to reassign the webzone processors.

```
TargetGuestDom# poolcfg -c 'create pset webzone-set(uint pset.min=32;uint pset.max=32)'
TargetGuestDom# poolcfg -c "create pool webzone-pool"
TargetGuestDom# poolcfg -c 'associate pool webzone-pool (pset webzone-set)'
TargetGuestDom# startcpu=16; endcpu=47
TargetGuestDom# i=$startcpu; while [ ${i} -le $endcpu ]; do set -x; poolcfg -c "transfer
 to pset webzone-set ( cpu ${i} )"; set +x; i= $((${i} + 1)); done
```

3. **Commit the changes.**

```
TargetGuestDom# pooladm -c
```

4. **Verify the processor configuration.**

```
root@TargetGuestDom# poolstat
                             pset
 id pool                 size used load
  4 webzone-pool           32 0.00 0.00
  0 pool_default           16 0.00 0.05
  3 dbzone-pool            32 0.00 0.00
```

▼ Reconfigure Host Names and IP Parameters

1. **List the zones.**

The zones are ready to be configured.

```
root@TargetGuestDom# zoneadm list -cv
```

```
  ID NAME             STATUS         PATH                         BRAND    IP
   0 global           running          /                          native   shared
   - webzone          installed        /rpool/webzone             native   shared
   - dbzone           installed        /zones/dbzone              native   shared
```

2. **Boot the non-global zones.**

   By issuing the following commands in separate terminal windows, you can watch for potential error messages for each zone, and also use the zone console to issue additional commands such as `sys-unconfig` (used in the next step).

   From a terminal window, log into the `dbzone` console:

   ```
   root@TargetGuestDom# zlogin -C dbzone
   ```

   From a second window, log into the `webzone` console:

   ```
   root@TargetGuestDom# zlogin -C webzone
   ```

   From a third terminal window, boot both zones:

   ```
   root@TargetGuestDom# zoneadm -z dbzone boot
   root@TargetGuestDom# zoneadm -z webzone boot
   ```

3. **Configure the zone's host names and IP parameters to conform with the target environment.**

   In this example, the following host names are identified to be ideal host names on the target system:

   `dbzone` → TargetDBzone.us.example.com

   `webzone` → TargetWebzone.us.example.com

   Use one of the following methods to change the hostnames and IP parameters on the target system:

   - **In each zone, update these files (if available) with the appropriate parameters for the new environment, and reboot the non-global zones.**

     - `/etc/domainname` – Specify the fully qualified domain name of the administrative domain to which the local host's network belongs.
     - `/etc/hosts` – Specify the zone's hostname and IP address
     - `/etc/hostname.`*interface* – (Only applies to exclusive-IP zones, which does not apply to this example) Specify the host name or IP address associated with the network interface.
     - `/etc/nodename` – Specify the host name of the local system.

- `/etc/resolve.conf` – Specify the domain name and name servers for this host.

```
root@TargetGuestDom# zoneadm -z dbzone reboot
root@TargetGuestDom# zoneadm -z webzone reboot
```

- **Alternatively, you can use `sys-unconfig` to reconfigure the non-global zones.**

  The `sys-unconfig` command resets a system's configuration so that it is ready to be reconfigured. The system's configuration consists of the hostname, network parameters, timezone, and root password. For more information, refer to the sys-unconfig man page.

  In this example, the command is performed in the `dbzone` console window.

```
root@TargetDBzone# sys-unconfig
```

  In this example, the command is performed in the `webzone` console window.

```
root@TargetWebzone# sys-unconfig
```

  From the target guest domain, boot the non-global zones.

  As the zones boot, watch the non-global zone console windows, and when prompted, enter the zone configuration information.

```
root@TargetGuestDom# zoneadm -z dbzone boot
root@TargetGuestDom# zoneadm -z webzone boot
```

4. **Verify the zone CPU assignments.**

```
root@TargetGuestDom# zlogin dbzone psrinfo -pv
The physical processor has 16 virtual processors (32-47)
  SPARC-S7 (chipid 1, clock 4267 MHz)
The physical processor has 16 virtual processors (48-63)
  SPARC-S7 (chipid 2, clock 4267 MHz)
root@TargetGuestDom# zlogin webzone psrinfo -pv
The physical processor has 16 virtual processors (0-15)
  SPARC-S7 (chipid 0, clock 4267 MHz)
The physical processor has 16 virtual processors (16-31)
  SPARC-S7 (chipid 1, clock 4267 MHz)
```

5. **Verify that the web server in `webzone` is accessible by a client.**

   A successful connection confirms that the `webzone` is deployed successfully on the target system.

Note - The same confirmation can be accomplished using a browser to display a web page from the web server.(`http://TargetWebzone.us.example.com`).

```
root@client-sys# /usr/sfw/bin/wget http://TargetWebzone.us.example.com
--2018-07-30 16:10:17--  http://TargetWebzone.us.example.com/
Resolving TargetWebzone.us.example.com (TargetWebzone.us.example.com)... 192.0.2.26

Connecting to TargetWebzone.us.example.com (TargetWebzone.us.example.com)|192.0.2.26|:
80... connected.

HTTP request sent, awaiting response... 200 OK
Length: 1456 (1.4K) [text/html]
Saving to: 'index.html.3'

index.html.3
100%[=======================================================================>]
 1.42K  --.-KB/s    in 0s

2018-07-30 16:10:17 (66.2 MB/s) - 'index.html.3' saved [1456/1456]
```

Transfer Data to the Database Zone

The `dbzone` on the source system has 2 zpools (`dbzone` and `dbzone_db_primary`) .The redo and archive logs are on a non-ZFS filesystem, and must be copied to the target `dbzone` before the `dbzone` database can be reconfigured:

- `dbzone` (zpool)
- `dbzone_db_primary` (zpool)
- `/logs/redologs` (UFS filesystem for Database Redologs)
- `/logs/archivelogs` (UFS filesystem for Database archivelogs)
- Target `dbzone` disks mapped to the data (ASM) contained in the Source `dbzone`

This procedure uses the following methods to move the application data to the target zones:

- The database logs are restored and copied to shared storage (`/ovas1`).
- A combination of `dd` and `gunzip` commands are used to get a consistent copy of the source disks (database application data) and remap the data to the target disks (same size) on the zones after they are attached to the target guest domain.

1. On the target guest domain, identify the disks.

```
root@TargetGuestDom# echo |format
AVAILABLE DISK SELECTIONS:
       0. c0d0 <SUN-DiskImage-16GB cyl 17064 alt 2 hd 96 sec 768>
          /virtual-devices@100/channel-devices@200/disk@0
```

```
       1. c0d1 <600GB disk for guest second root dis cyl 17064 alt 2 hd 96 sec 768>
          /virtual-devices@100/channel-devices@200/disk@1
       2. c0d2 <SUN-ZFS Storage 7355-1.0-300.00GB>
          /virtual-devices@100/channel-devices@200/disk@2
       3. c0d3 <SUN-ZFS Storage 7355-1.0-300.00GB>
          /virtual-devices@100/channel-devices@200/disk@3
       4. c0d4 <SUN-ZFS Storage 7355-1.0-150.00GB>
          /virtual-devices@100/channel-devices@200/disk@4
       5. c0d5 <SUN-ZFS Storage 7355-1.0-150.00GB>
          /virtual-devices@100/channel-devices@200/disk@5
       6. c0d6 <SUN-ZFS Storage 7355-1.0-200.00GB>
          /virtual-devices@100/channel-devices@200/disk@6
       7. c0d7 <SUN-ZFS Storage 7355-1.0-200.00GB>
          /virtual-devices@100/channel-devices@200/disk@7
       8. c0d8 <SUN-ZFSStorage7355-1.0 cyl 8124 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@8
       9. c0d9 <SUN-ZFSStorage7355-1.0 cyl 8124 alt 2 hd 254 sec 254>
          /virtual-devices@100/channel-devices@200/disk@9
      10. c0d10 <SUN-ZFSStorage7355-1.0 cyl 8190 alt 2 hd 8 sec 32>
          /virtual-devices@100/channel-devices@200/disk@a
      11. c0d11 <SUN-ZFSStorage7355-1.0 cyl 8190 alt 2 hd 8 sec 32>
          /virtual-devices@100/channel-devices@200/disk@b
Specify disk (enter its number): Specify disk (enter its number):
```

2. Restore the redo and archive logs needed for the database startup on `dbzone`.

To copy the contents of the Database and ASM data (diskgroup, and so on) these steps are carried out on the `dbzone` zone.

a. Mount the shared storage from the target control domain.

The shared mount point from the target control domain is mounted on `dbzone` to provide access to the application data disk contents.

```
root@TargetDBzone# mkdir /ovas1
root@TargetDBzone# mount -F nfs TargetControlDom:/ovas1 /ovas1
```

b. Copy the redo logs to the target zone.

```
root@TargetDBzone# cd /logs/redologs/
root@TargetDBzone# gunzip -c /ovas1/redo.ufsdump.gz | ufsrestore xvf -
Verify volume and initialize maps
Dump   date: Mon Jul 30 11:09:29 2018
Dumped from: the epoch
Level 0 dump of /zones/dbzone/root/logs/redologs on SourceGlobal:/dev/md/dsk/d20
Label: none
Extract directories from tape
Initialize symbol table.
```

```
Extract requested files
extract file ./redo04.log
extract file ./redo05.log
extract file ./redo06.log
Add links
Set directory mode, owner, and times.
set owner/mode for '.'? [yn] y

root@TargetDBzone# ls -rlth redo*
-rw-r-----   1 oracle1  dba         200M Jul 30 11:00 redo05.log
-rw-r-----   1 oracle1  dba         200M Jul 30 11:02 redo06.log
-rw-r-----   1 oracle1  dba         200M Jul 30 11:09 redo04.log
```

c. Copy the archive logs to the target zone.

```
root@TargetDBzone# cd /logs/archivelogs/
root@TargetDBzone# gunzip -c  /ovas1/archive.ufsdump.gz | ufsrestore xvf -
Verify volume and initialize maps
Dump   date: Mon Jul 30 11:09:29 2018
Dumped from: the epoch
Level 0 dump of /zones/dbzone/root/logs/archivelogs on SourceGlobal:/dev/md/dsk/d30
Label: none
Extract directories from tape
Initialize symbol table.
Make node ./ORCL9
Make node ./ORCL9/archivelog
Make node ./ORCL9/archivelog/2018_07_17
Make node ./ORCL9/archivelog/2018_07_26
Make node ./ORCL9/archivelog/2018_07_27
Make node ./ORCL9/archivelog/2018_07_28
Make node ./ORCL9/archivelog/2018_07_29
Extract requested files
extract file ./ORCL9/archivelog/2018_07_17/o1_mf_1_50_fnvzwkj8_.arc
extract file ./ORCL9/archivelog/2018_07_26/o1_mf_1_51_foo2omnh_.arc
extract file ./ORCL9/archivelog/2018_07_26/o1_mf_1_52_foo9vbdk_.arc
extract file ./ORCL9/archivelog/2018_07_26/o1_mf_1_53_foobh0oh_.arc
extract file ./ORCL9/archivelog/2018_07_27/o1_mf_1_54_foojrfps_.arc
extract file ./ORCL9/archivelog/2018_07_27/o1_mf_1_55_foqy30ls_.arc
extract file ./ORCL9/archivelog/2018_07_28/o1_mf_1_56_forhorlz_.arc
extract file ./ORCL9/archivelog/2018_07_28/o1_mf_1_57_fosdvstz_.arc
extract file ./ORCL9/archivelog/2018_07_28/o1_mf_1_58_fotb98gv_.arc
extract file ./ORCL9/archivelog/2018_07_29/o1_mf_1_59_fovgm1cx_.arc
extract file ./ORCL9/archivelog/2018_07_29/o1_mf_1_60_fow18vt2_.arc
extract file ./ORCL9/archivelog/2018_07_29/o1_mf_1_61_fowydo5d_.arc
Add links
Set directory mode, owner, and times.
set owner/mode for '.'? [yn] y
```

```
root@TargetDBzone# ls -lrtRh ORCL9
ORCL9:
total 16
drwxr-x---   7 oracle1  dba          452 Jul 29 06:00 archivelog

ORCL9/archivelog:
total 80
drwxr-x---   2 oracle1  dba          198 Jul 17 07:39 2018_07_17
drwxr-x---   2 oracle1  dba          360 Jul 26 22:13 2018_07_26
drwxr-x---   2 oracle1  dba          279 Jul 27 22:00 2018_07_27
drwxr-x---   2 oracle1  dba          360 Jul 28 19:40 2018_07_28
drwxr-x---   2 oracle1  dba          360 Jul 29 19:36 2018_07_29

ORCL9/archivelog/2018_07_17:
total 4112
-rw-r-----   1 oracle1  dba         2.0M Jul 17 07:39 o1_mf_1_50_fnvzwkj8_.arc

ORCL9/archivelog/2018_07_26:
total 445936
-rw-r-----   1 oracle1  dba          73M Jul 26 20:00 o1_mf_1_51_foo2omnh_.arc
-rw-r-----   1 oracle1  dba          74M Jul 26 22:02 o1_mf_1_52_foo9vbdk_.arc
-rw-r-----   1 oracle1  dba          71M Jul 26 22:13 o1_mf_1_53_foobh0oh_.arc

ORCL9/archivelog/2018_07_27:
total 299728
-rw-r-----   1 oracle1  dba          72M Jul 27 00:00 o1_mf_1_54_foojrfps_.arc
-rw-r-----   1 oracle1  dba          75M Jul 27 22:00 o1_mf_1_55_foqy30ls_.arc

ORCL9/archivelog/2018_07_28:
total 443696
-rw-r-----   1 oracle1  dba          73M Jul 28 03:00 o1_mf_1_56_forhorlz_.arc
-rw-r-----   1 oracle1  dba          72M Jul 28 11:18 o1_mf_1_57_fosdvstz_.arc
-rw-r-----   1 oracle1  dba          72M Jul 28 19:40 o1_mf_1_58_fotb98gv_.arc

ORCL9/archivelog/2018_07_29:
total 442768
-rw-r-----   1 oracle1  dba          73M Jul 29 06:00 o1_mf_1_59_fovgm1cx_.arc
-rw-r-----   1 oracle1  dba          71M Jul 29 11:19 o1_mf_1_60_fow18vt2_.arc
-rw-r-----   1 oracle1  dba          72M Jul 29 19:36 o1_mf_1_61_fowydo5d_.arc
```

3. Restore ASM data.

```
root@TargetDBzone# gunzip -c /ovas1/asm1.img.gz | dd of=/dev/rdsk/c0d6s0 bs=104857600
root@TargetDBzone# gunzip -c /ovas1/asm2.img.gz | dd of=/dev/rdsk/c0d7s0 bs=104857600
```

4. Change the ownership and the permissions of the ASM disks.

The ownership of the ASM disks should owned by `oracle1:dba` and have permissions set to 660.

```
root@TargetDBzone# chown oracle1:dba /dev/rdsk/c0d6s0 /dev/rdsk/c0d7s0
root@TargetDBzone# chmod 660 /dev/rdsk/c0d6s0 /dev/rdsk/c0d7s0
```

5. **Check the networking topology of the `dbzone`.**

```
root@TargetDBzone# ifconfig -a
lo0:1: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
        inet 127.0.0.1 netmask ff000000
vnet0:1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 2
        inet 192.0.2.52 netmask fffffe00 broadcast 192.0.2.254

TargetDBzone# netstat -rn

Routing Table: IPv4
  Destination           Gateway       Flags  Ref     Use     Interface
-------------------- --------------- ----- ----- ---------- ---------
default              192.0.2.1         UG        1        225
192.0.2.0            192.0.2.52        U         1          2 vnet0:1
224.0.0.0            192.0.2.52        U         1          0 vnet0:1
127.0.0.1            127.0.0.1         UH        8      49807 lo0:1
```

▼ Reconfigure the Database

The Oracle Database 12.1.0.2 environment has moved to a new hostname `TargetDBzone`, so the database requires some reconfiguration.

1. **Perform all steps except for the last step listed in:**

   Configure Guest Domain Database Components (a chapter in the *Lift and Shift Guide - Moving Oracle Solaris 10 Guest Domains to SPARC Servers Running Oracle Solaris 11*), then return to Step 2 in this procedure.

   The Oracle Database environment moved to a domain with a new hostname, and therefore requires some reconfiguration. The referenced procedure describes the reconfiguration process. The steps were derived and modified from the MOS document titled How to Reconfigure Oracle Restart on 12c / 12.1 (Doc ID 1570358.1).

2. **Recreate the non-default database service.**

```
TargetDBzone$ srvctl add service -d orcl9 -pdb pdborcl -s mydb
TargetDBzone$ srvctl status service -d orcl9 -s mydb
Service mydb is not running.

TargetDBzone$ srvctl start service -d orcl9 -s mydb
TargetDBzone$ srvctl status service -d orcl9 -s mydb
```

```
Service mydb is running
```

3. **On `dbzone`, check the status of the database.**

```
TargetDBzone$ crsctl status resource -t
--------------------------------------------------------------------------------
Name           Target  State        Server                   State details
--------------------------------------------------------------------------------
Local Resources
--------------------------------------------------------------------------------
ora.DATA.dg
               ONLINE  ONLINE       TargetDBzone STABLE
ora.LISTENER.lsnr
               ONLINE  ONLINE       TargetDBzone STABLE
ora.asm
               ONLINE  ONLINE       TargetDBzone Started,STABLE
ora.ons
               OFFLINE OFFLINE      TargetDBzone STABLE
--------------------------------------------------------------------------------
Cluster Resources
--------------------------------------------------------------------------------
ora.cssd
      1        ONLINE  ONLINE       TargetDBzone STABLE
ora.diskmon
      1        OFFLINE OFFLINE                               STABLE
ora.evmd
      1        ONLINE  ONLINE       TargetDBzone STABLE
ora.orcl9.db
      1        ONLINE  ONLINE       TargetDBzone Open,STABLE
ora.orcl9.mydb.svc
      1        ONLINE  ONLINE       TargetDBzone STABLE
--------------------------------------------------------------------------------
```

4. **Check the database listener to ensure that the database services are successfully migrated.**

```
TargetDBzone$ lsnrctl status

LSNRCTL for Solaris: Version 12.1.0.2.0 - Production on 17-AUG-2018 09:01:06
Copyright (c) 1991, 2014, Oracle.  All rights reserved.
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=TargetDBzone.us.example.com)
(PORT=1521)))
STATUS of the LISTENER
------------------------
Alias                     LISTENER
Version                   TNSLSNR for Solaris: Version 12.1.0.2.0 - Production
Start Date                17-AUG-2018 09:00:33
Uptime                    0 days 0 hr. 0 min. 32 sec
```

```
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Listener Parameter File   /u01/app/oracle/product/12.1.0/grid/network/admin/listener.ora
Listener Log File         /u01/app/oracle/diag/tnslsnr/TargetDBzone/listener/alert/
log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=TargetDBzone.us.example.com)(PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=TargetDBzone.us.example.com)(PORT=5500))
(Security=(my_wallet_directory=/u01/app/oracle/product/12.1.0/dbhome_1/admin/orcl9/
xdb_wallet))(Presentation=HTTP)(Session=RAW))
Services Summary...
Service "mydb.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
Service "orcl9.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
Service "orcl9XDB.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
Service "pdborcl.us.example.com" has 1 instance(s).
  Instance "orcl9", status READY, has 1 handler(s) for this service...
The command completed successfully
-bash-3.2$
```

5. **From a database client system, check that the database files are in expected locations.**

```
root@client-sys# sqlplus system/welcome1@//TargetDBzone.us.example.com:1521/mydb.us.
example.com

SQL> select name from v$datafile union select member from v$logfile union select name
 from v$archived_log;
+DATA/ORCL9/DATAFILE/system.258.981701267
+DATA/ORCL9/DATAFILE/sysaux.257.981701213
...............
+DATA/ORCL9/71335771C4785E92E054000B5DDC023F/DATAFILE/soe.278.981702033
.....
/logs/redologs/redo04.log
/logs/redologs/redo05.log
...........
......
/logs/archivelogs/ORCL9/archivelog/2018_07_17/o1_mf_1_50_fnvzwkj8_.arc
......
```

6. **Check the application schema (SOE) to verify that the expected data was migrated.**

The count should match the count from “Lift the Source Environment to the Shared Storage” on page 49.

```
SQL> select count(*) from soe.orders ;

  COUNT(*)
----------
  1429790
```

7. **Enable the CRS stack.**

   The CRS stack was disabled in the beginning of the migration. Now that the data has been lifted and shifted to the target, the CRS stack must be enabled.

   Run this command as the application user (`oracle1`):

```
TargetDBzone$ crsctl enable has
```

▼ Perform Assurance Tests and Release the Workloads into Production

1. **Use Oracle Enterprise Manager Database Express to check the status of the database.**

   Note – For more information on Enterprise Manager Database Express, refer to this web page:
   https://www.oracle.com/technetwork/database/manageability/emx-intro-1965965.html

In this example, the status and performance characteristics show that the database was successfully migrated to the target system.

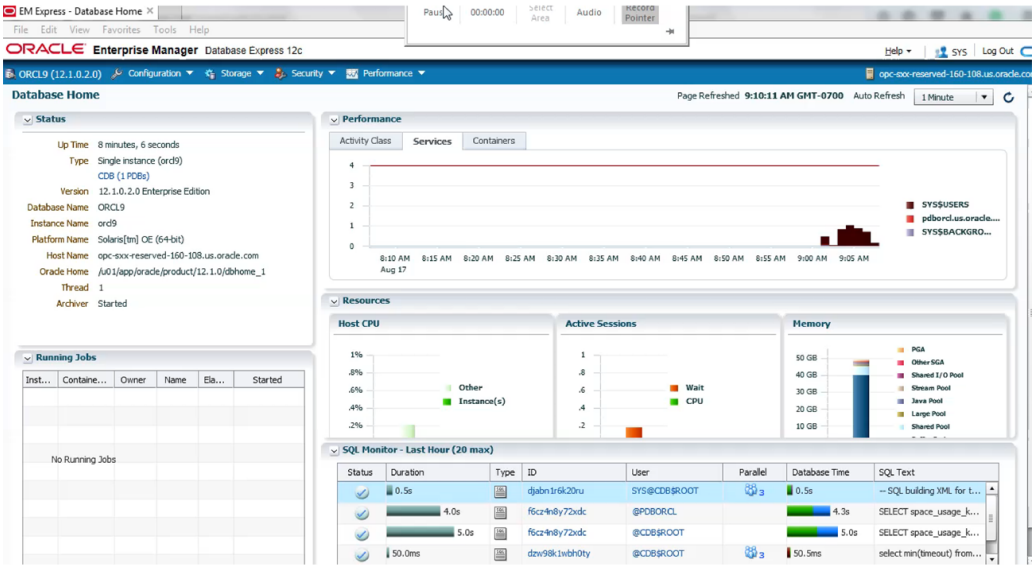

2. **Use the assurance tests for your environment to ensure the migrated zones are operable.**

   Performance data can be compared to the data collected on the source system. In particular, compare the CPU utilization. In some cases, the new system provides improved compute performance for the workloads in the target guest domain, and might result in significantly lower CPU utilization than the source system. In this case, consider assigning fewer CPU resources to the guest domain. Start by reducing the guest domain CPU processor sets which are used in the non-global zones. Then reduce the number of cores allocated to the guest domain. Using fewer CPU cores can reduce software licensing costs.

3. **Release the migrated workloads for use in your environment.**

Index

M

N

O

P

R

S

T

U

Z