

**Lift and Shift Guide - Migrating
Workloads from Oracle Solaris 10 SPARC
Systems to Oracle Solaris 10 Branded
Zones**

ORACLE

Part No: F10607-03
February 2020

Part No: F10607-03

Copyright © 2020, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

Access to Oracle Support

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

Référence: F10607-03

Copyright © 2020, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

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

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

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

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

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

Accès aux services de support Oracle

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

Contents

Using This Documentation	7
Product Documentation Library	7
Feedback	7
Understanding the Lift and Shift Process	9
Migration Methods Based on the OS Version	9
Case 1 (Preferred) – The Source System is Running Oracle Solaris 10 Update 9 (or later)	10
Case 2 – The Source System is Running Oracle Solaris 10 Update 8 (or Earlier)	10
Lift and Shift Overview	11
Lift and Shift Process	13
Requirements	15
Example Configuration and Prompts	16
Preparing the Lift and Shift Systems	19
▼ Check the Source System's Suitability for Migration	19
▼ Review the Source System Configuration	21
▼ Prepare the Shared Storage	29
▼ Prepare the Target System	32
▼ Configure an Oracle Solaris 10 Branded Zone on the Target System	40
Lifting and Shifting the Workloads	45
▼ Lift the Workloads to the Shared Storage	45
▼ Shift the Workloads to the Target	47
Verifying the New Zone and Workloads	51
▼ Verify the Solaris 10 Branded Zone Configuration	51

▼ Start and Verify the Migrated Database	54
▼ Perform Assurance Tests and Release the Workloads into Production	57
Index	59

Using This Documentation

- **Overview** – This document provides instructions for lifting and shifting workloads running on Oracle Solaris 10 with UFS or ZFS root file systems to Oracle Solaris 10 branded zone in Oracle Solaris 11.3 (or later) running on more modern sun4v hardware.
- **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:

- “Migration Methods Based on the OS Version” on page 9
- “Lift and Shift Overview” on page 11
- “Requirements” on page 15
- “Example Configuration and Prompts” on page 16

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/

Migration Methods Based on the OS Version

This document provides instructions for lifting and shifting workloads running on Oracle Solaris 10 with UFS or ZFS root file systems to Oracle Solaris 10 branded zone in Oracle Solaris 11.3 (or later) running on more modern sun4v hardware. The specific methods that you use to perform this lift and shift scenario differ based on the version of Oracle Solaris 10 on the source system.

To identify the version of Oracle Solaris, view the `/etc/release` file. Examples:

```
# cat /etc/release
Solaris 10 10/09 s10s_u8wos_08a SPARC
Copyright 2009 Sun Microsystems, Inc. All Rights Reserved.
Use is subject to license terms.
Assembled 16 September 2009

# cat /etc/release
Oracle Solaris 10 1/13 s10s_u11wos_24a SPARC
Copyright (c) 1983, 2013, Oracle and/or its affiliates. All rights reserved.
Assembled 17 January 2013
```

The following table can help you identify the update version of Oracle Solaris 10 based on the release date.

Release Date	Oracle Solaris 10	Release Date	Oracle Solaris 10
03/05	n/a	10/08	Update 6
01/06	Update 1	05/09	Update 7
06/06	Update 2	10/09	Update 8
11/06	Update 3	09/10	Update 9
08/07	Update 4	08/11	Update 10
05/08	Update 5	01/13	Update 11

Case 1 (Preferred) – The Source System is Running Oracle Solaris 10 Update 9 (or later)

Follow the instructions in this document.

Case 2 – The Source System is Running Oracle Solaris 10 Update 8 (or Earlier)

Choose from one of these methods:

Method 1 – Upgrade the source system to Oracle Solaris 10 Update 9 (or later). Then follow the instructions in this document.

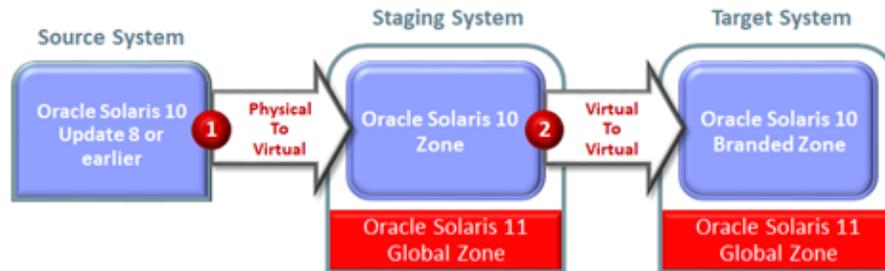
For information about upgrading, select the appropriate Oracle Solaris 10 documentation library at https://docs.oracle.com/cd/F24622_01

Method 2 – Follow the instructions in this document and ensure that you use the `zoneadm install -P` option as described in “[Shift the Workloads to the Target](#)” on page 47.

The `zoneadm install -P` option automatically applies required patches to migrated branded zone on the target system.

By default, the `zonep2vchk` utility that is used as a pre-migration checker does not exist in Oracle Solaris 10 Update 8 or earlier. However, if you have access to any higher versions of Oracle Solaris OS, you can copy the `/usr/sbin/zonep2vchk` utility to the source system and run the utility as described in this guide.

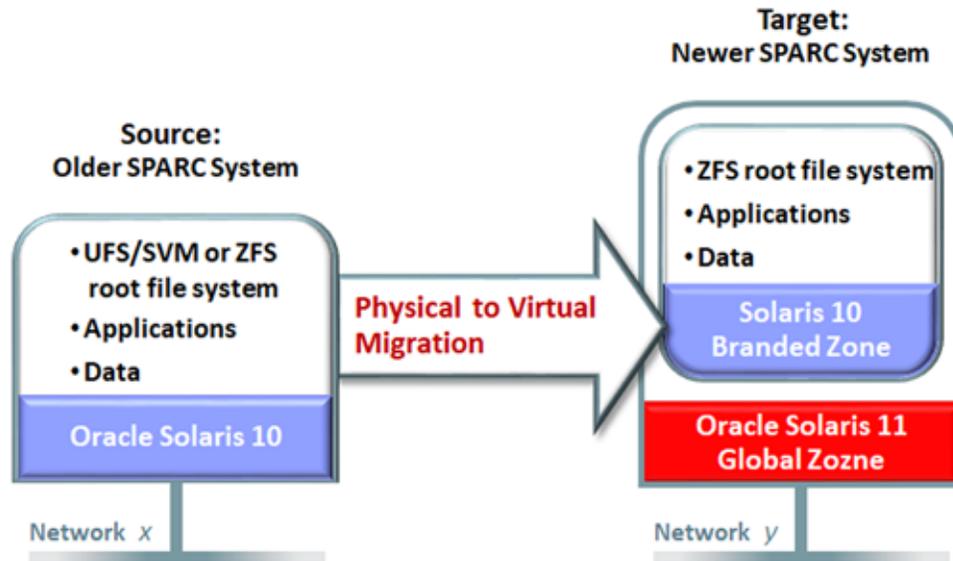
Method 3 – The migration is performed in two migration phases using a staging system.



1. Migrate the source environment to an Oracle Solaris 10 zone on a staging system that is running Oracle Solaris 11 in the global zone. Refer to guide titled [Lift and Shift Guide - Moving Oracle Solaris 10 Guest Domains to SPARC Servers Running Oracle Solaris 11](#)
2. Migrate the Oracle Solaris 10 zone to an Oracle Solaris 10 branded zone on the target system (running Oracle Solaris 11.3 or later) using the instructions in this document.

Lift and Shift Overview

This document provides instructions for performing a physical to virtual lift and shift. The techniques are specifically aimed at migrating workloads running on Oracle Solaris 10 with UFS or ZFS root file systems on SPARC systems to Oracle Solaris 10 branded zones in an Oracle Solaris 11.3 (or later) global zone.

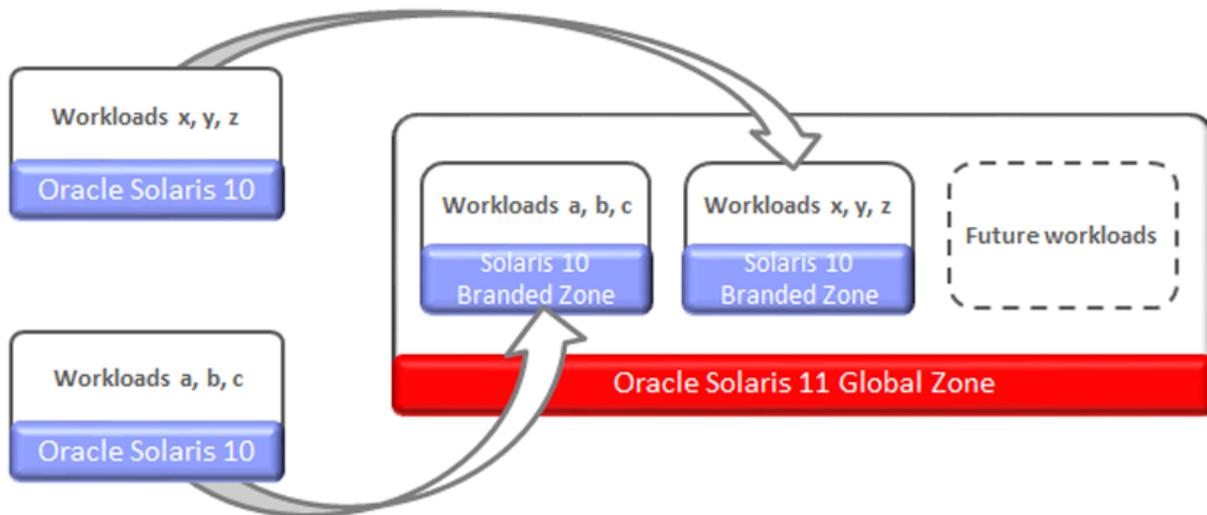


Oracle Solaris 10 branded zones, also known as `solaris10` branded non-global zones, run Oracle Solaris 10 applications in a virtualized Oracle Solaris 10 environment on top of the Oracle Solaris 11 operating system. Applications run unmodified in the secure environment provided by the non-global zone, and can take advantage of the enhancements made to the kernel and utilize some of the innovative technologies that are only available in Oracle Solaris 11 releases.

For additional information about Oracle Solaris zones, refer to the zone documentation in the Oracle Solaris Information Library for your specific release:

- [Oracle Solaris 11.4 Information Library](#)
- [Oracle Solaris 11.3 Information Library](#)
- [Access to other Oracle Solaris Libraries](#)

With a higher number of compute resources on newer SPARC systems, a single global zone can host a higher number of non-global zones, consolidating compute workloads. This lift and shift procedure can be used for each Oracle Solaris 10 single system that you plan to move to the newer SPARC system.

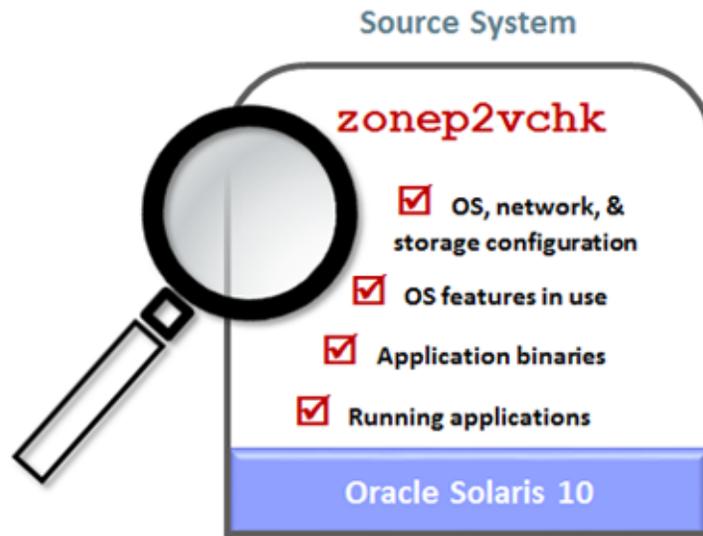


Lift and Shift Process

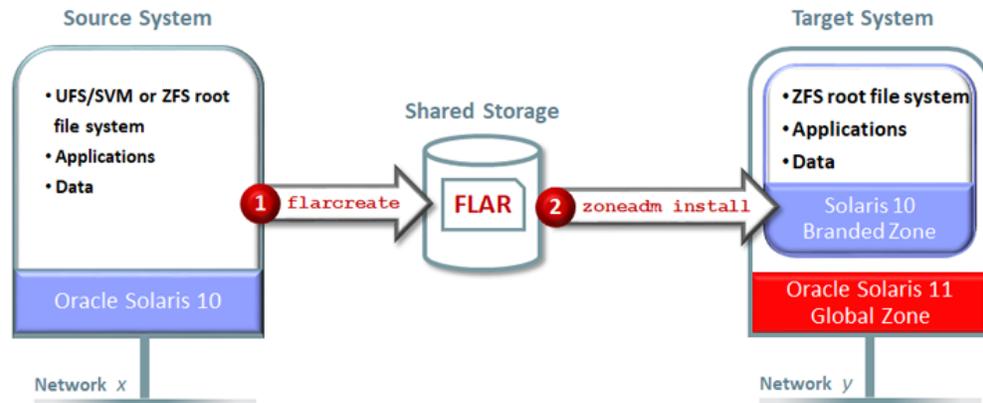
The lift and shift process uses a variety of Oracle Solaris utilities to migrate the entire source configuration, including all associated applications and data, to a `solaris10` branded non-global zone on target system's Oracle Solaris 11 global zone. To accomplish this migration, the following activities are performed, although not necessarily in this exact order. The procedures in subsequent chapters order the activities to minimize the amount of time when services are unavailable.

Note - 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 in this case), it is a best practice to shutdown services during the lift and shift activities.

- Analyze the source system to ensure migration suitability** –(Note - This command is only available in Oracle Solaris 10 Update 9 or later.) One way to analyze the source system is to run the Oracle Solaris `zonep2vchk` utility. This utility evaluates a global zone or physical instance for potential migration issues. The utility can detect issues with the OS, network, and storage configuration, identify some OS features, application binaries, and applications that might not be supported in a non-global zone,



- **Assess the source system resources and workloads** – 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, network, and workload resources.
- **Select a 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 shared storage** – Storage on the target system is exported as a network file system and mounted on the source system, for the efficiency of the lift and shift process.
- **Prepare the target system** – CPU, memory, networking and storage resources are configured to accommodate the incoming Oracle Solaris 10 environment.
- **Configure a solaris10 branded non-global zone on the target system** – A new solaris10 branded non-global zone is configured, but not yet installed, on the target's global zone.
- **Lift** – The `flarcreate` command is used to capture the source system's file systems, including the OS, applications, and data, and compressing those items into a flash archive (FLAR) on shared storage.
- **Shift** – The `zoneadm install` command uses the FLAR to restore the source OS root file system to a ZFS root file system in the Solaris 10 branded zone.



- **Reconfigure** – The shifted environment might require post-migration configuration adjustments to function in the new environment. For example, if the network configuration is carried over from the source environment, the network parameters might need to be changed to function in the new environment.
- **Verify functionality** – Verify the functionality of the workloads. If possible, take performance measurements and adjust resources to run optimally in the new environment.

Requirements

The lift and shift process described in this document has the following requirements:

Source system

- Must be a bare metal or logical domain running the Oracle Solaris 10 OS. The example in this document uses a bare metal system running the Oracle Solaris 10 1/13 OS.
- Must have these patches installed, at minimum: 119254-75, 119534-24, and 140914-02 or 148027-03 (The latter patch obsoletes the former patch. Either patch is sufficient for this migration.) Instructions in this document describe how to check and install the patches.

These patches ensure that the package and patch tools work properly within the solaris10 branded zone.

Target system

Must be running Oracle Solaris 11.3 or later. The examples in this document assume that the target system's OS was freshly installed and non-global zones have not yet been created.

Shared storage

You must provide sufficient shared storage space to temporarily store data that is used to perform the lift and shift. The storage must be accessible to the source and target systems. For instructions about determining the space requirements and preparing the storage, see [“Prepare the Shared Storage” on page 29](#).

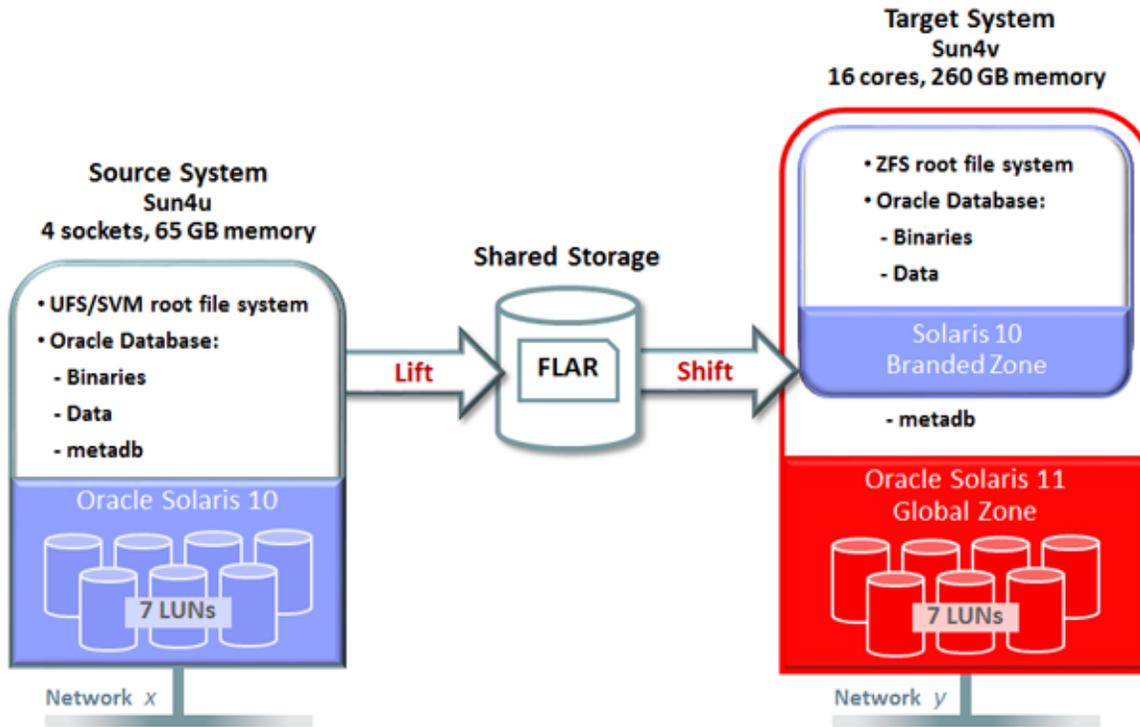
Example Configuration and Prompts

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.

The following systems are used in the examples:

- **Source system** – A SPARC Enterprise M9000 (sun4u) system running the Oracle Solaris 10 OS and Oracle Database 12.1.0.2.
- **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 a Solaris 10 branded zone that hosts the OS and workloads from the source system.

The following diagram shows the source and target configurations that are used in the examples in this guide.



In the screen output examples, the command line prompt indicates on which system (target or source), and which zone (global or non-global) a command is executed. This table lists the prompts.

TABLE 1 System Identification Prompts

System	Component	Host Name in the Prompt
Source System	Oracle Solaris 10 physical instance	SourceSystem
Target System	Oracle Solaris 11 global zone	TargetGlobal
	Solaris 10 branded non-global zone	TargetS10bz
Other Systems	A client system on the network for testing client access.	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:

- [“Check the Source System's Suitability for Migration” on page 19](#)
- [“Review the Source System Configuration” on page 21](#)
- [“Prepare the Shared Storage” on page 29](#)
- [“Prepare the Target System” on page 32](#)
- [“Configure an Oracle Solaris 10 Branded Zone on the Target System” on page 40](#)

▼ Check the Source System's Suitability for Migration

1. **Log into the source global zone as superuser.**
2. **Run the `zonep2vchk` command.**

```
root@SourceSystem# zonep2vchk
--Executing Version: 5.10.1.1
- Source System: SourceSystem
  Solaris Version: Oracle Solaris 10 1/13 s10s_u11wos_24a SPARC
  Solaris Kernel: 5.10 Generic_150400-61
  Platform: sun4u SUNW,SPARC-Enterprise
- Target System:
  Solaris Version: Solaris 10
  Zone Brand: native (default)
  IP type: shared
--Executing basic checks
- The following /etc/system tunables exist. These tunables will not
  function inside a zone. The /etc/system tunable may be transferred to
  the target global zone, but it will affect the entire system,
  including all zones and the global zone. If there is an
  alternate tunable that can be configured from within the zone,
  this tunable is described:
  rootdev:/pseudo/md@0:0,10,blk
```

```
zonep2vchk has no information on tunable

set md:mirrored_root_flag = 1
zonep2vchk has no information on tunable

- The system is configured with the following ISCSI initiators. A zone
cannot access ISCSI targets. ISCSI targets must be discovered and
configured from the global zone. See iscsiadm(1M) for details:
    iqn.1986-03.com.sun:02:3e10566b-4e39-e82f-b866-ec2d56d14278
- The following SMF services will not work in a zone:
    svc:/network/iscsi/initiator:default
    svc:/system/iscsitgt:default
- The following SMF services require ip-type "exclusive" to work in
a zone. If they are needed to support communication after migrating
to a shared-IP zone, configure them in the destination system's global
zone instead:
    svc:/network/ipsec/ipsecalgs:default
    svc:/network/ipsec/policy:default
    svc:/network/routing-setup:default
- The system is configured with SVM metadevices.
A zone cannot configure SVM metadevices, but a zone can be configured
to use existing metadevices. See metastat(1M) for details
- If needed, the following non-standard vfstab entries will impact
the zone configuration:
    Device
    /dev/md/dsk/d30                /dev/md/rdisk/d30 /u01 ufs 2 yes
forcedirectio,logging /dev/md/dsk/d30
    /dev/md/dsk/d40                /dev/md/rdisk/d40 /oradata ufs 2 yes
forcedirectio,logging /dev/md/dsk/d40
- When migrating to an exclusive-IP zone, the target system must have an
available physical interface for each of the following source system
interfaces:
    bge0
    bge2
- When migrating to an exclusive-IP zone, interface name changes may
impact the following configuration files:
    /etc/hostname.bge0
    /etc/hostname.bge2
- When migrating to a shared-IP zone, the following network features must
be configured in the global zone if needed to support communication.
They will not be configured automatically during migration. Nothing
needs to be done for these features when migrating to an exclusive-IP
zone where the configuration inside the migrated system image will be
used:
    IP Multipath group ipmp0
Basic checks complete. Issue(s) detected: 16
--Total issue(s) detected: 16
```

3. Examine this report line by line before proceeding with the migration.

In this example, there are no critical issues, and the few reported issues are addressed later in [“Prepare the Target System” on page 32](#) and [“Verify the Solaris 10 Branded Zone Configuration” on page 51](#).

▼ Review the Source System Configuration

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

2. (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@SourceSystem# script /tmp/source_zones_output.txt
```

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

3. Check the version of the operating system.

This lift and shift process is specifically for a source system running Oracle Solaris 10.

```
root@SourceSystem# cat /etc/release
Oracle Solaris 10 1/13 s10s_u11wos_24a SPARC
Copyright (c) 1983, 2013, Oracle and/or its affiliates. All rights reserved.
Assembled 17 January 2013
```

4. List the source system architecture and the kernel patch version.

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

5. Ensure that the required patches are installed.

a. Determine if the patches are installed and at or above the minimum versions.

Minimum versions:

119254-75

119534-24

140914-02 or **148027-03** (The latter patch obsoletes the former patch. Either patch is sufficient for this migration.)

```
root@SourceSystem# showrev -p | egrep "Patch: 119254|Patch: 119534|Patch: 140914|
Patch: 148027"
Patch: 119254-88 Obsoletes: 119015-03 Requires: 121133-02 Incompatibles: Packages:
  SUNWinstall-patch-utils-root, SUNWpkgcmds, SUNWpkgcmds, SUNWswmt
Patch: 119534-33 Obsoletes: Requires: 119252-18, 120199-09, 126677-02
  Incompatibles: Packages: SUNWinst
Patch: 148027-03 Obsoletes: 121002-04, 126316-01, 126651-02, 127920-01, 127922-04,
  128330-02, 137088-01, 138275-01, 138621-02, 138623-05, 140914-02,
  142009-01, 142336-01, 143588-01, 144300-01, 144876-01, 146578-06 Requires: 118833-
  36, 120011-14, 127127-11, 137137-09, 139555-08, 142909-17
```

- b. For any missing or below-version patches, download and install the latest versions offered from My Oracle Support (<http://support.oracle.com>)**

For patch details, refer to the patch README.

Use the patchadd command to install each patch. For example:

```
root@SourceSystem# patchadd ./119254-93
```

- 6. List the amount of memory.**

```
root@SourceSystem# prttdiag|grep "Memory size"
Memory size: 65536 Megabytes
```

- 7. List the source system processors.**

In this example, the source system has 4 sockets.

```
root@SourceSystem# psrinfo -pv
The physical processor has 8 virtual processors (0-7)
  SPARC64-VII+ (portid 1024 impl 0x7 ver 0xc0 clock 3000 MHz)
The physical processor has 8 virtual processors (40-47)
  SPARC64-VII+ (portid 1064 impl 0x7 ver 0xc0 clock 3000 MHz)
The physical processor has 8 virtual processors (80-87)
  SPARC64-VII+ (portid 1104 impl 0x7 ver 0xc0 clock 3000 MHz)
The physical processor has 8 virtual processors (120-127)
  SPARC64-VII+ (portid 1144 impl 0x7 ver 0xc0 clock 3000 MHz)
```

- 8. Display network configuration information.**

```
root@SourceSystem# 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 198.51.100.69 netmask ffffffff broadcast 198.51.100.255
    groupname ipmp0
    ether 0:b:5d:dc:2:40
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:2:40
sppp0: flags=10010008d1<UP,POINTOPOINT,RUNNING,NOARP,MULTICAST,IPv4,FIXEDMTU> mtu 1500
    index 4
    inet 203.0.113.3 --> 203.0.113.1 netmask ffffffff
    ether 0:0:0:0:0:0

```

```

root@SourceSystem# netstat -rn
Routing Table: IPv4

```

Destination	Gateway	Flags	Ref	Use	Interface
default	198.51.100.1	UG	1	1135000	
198.51.100.0	198.51.100.69	U	1	13	bge0
203.0.113.1	203.0.113.3	UH	1	1	sppp0
224.0.0.0	198.51.100.69	U	1	0	bge0
127.0.0.1	127.0.0.1	UH	5	4116	lo0

```

root@SourceSystem# ls -l /etc/hostname*
-rw-r--r--  1 root   root      46 Apr 19 16:18 /etc/hostname.bge0
-rw-r--r--  1 root   root      15 Apr 19 16:18 /etc/hostname.bge2

```

```

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

```

```

root@SourceSystem# cat /etc/hostname.bge2
group ipmp0 up

```

```

root@SourceSystem# cat /etc/inet/hosts
# Internet host table
::1      localhost
127.0.0.1 localhost
198.51.100.69 SourceSystem      loghost

```

9. Identify the source system disks and what is stored in them.

While performing the following substeps, make note of the amount of storage used by the various software components. In the last substep, storage values are tallied to determine the amount of storage space that is needed in shared storage and on the target system.

a. Use the format utility to list the disks.

The list of disks provides information about the disks on the source system.

In this example, disks number 0, 1, and 2 are physical internal disks. Disks 3 - 6 (where each disk is identified by a WWN) are LUNS.

```
root@SourceSystem# echo | format

Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <FUJITSU-MAY2073RC-3701 cyl 14087 alt 2 hd 24 sec 424>
     /pci@0,600000/pci@0/scsi@1/sd@0,0
  1. c0t1d0 <SUN600G 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. c2t600144F0E635D8C700005AC804B80019d0 <SUN-ZFSSStorage7420-1.0 cyl 4873 alt
     2 hd 254 sec 254>
     /scsi_vhci/ssd@g600144f0e635d8c700005ac804b80019
  4. c2t600144F0E635D8C700005AC804D4001Ad0 <SUN-ZFSSStorage7420-1.0 cyl 4873 alt
     2 hd 254 sec 254>
     /scsi_vhci/ssd@g600144f0e635d8c700005ac804d4001a
  5. c2t600144F0E635D8C700005AC804F4001Bd0 <SUN-ZFSSStorage7420-1.0 cyl 6499 alt
     2 hd 254 sec 254>
     /scsi_vhci/ssd@g600144f0e635d8c700005ac804f4001b
  6. c2t600144F0E635D8C700005AC8050C001Cd0 <SUN-ZFSSStorage7420-1.0 cyl 6499 alt
     2 hd 254 sec 254>
     /scsi_vhci/ssd@g600144f0e635d8c700005ac8050c001c
Specify disk (enter its number): Specify disk (enter its number):
```

b. List the root file system.

In this example, the root file system is on an SVM root encapsulated mirror.

```
root@SourceSystem# df -h /
Filesystem          size  used  avail capacity  Mounted on
/dev/md/dsk/d10     487G  4.9G   477G     2%      /
root@SourceSystem# metastat -p d10
d10 -m d11 d12 1
d11 1 1 c0t1d0s0
d12 1 1 c1t0d0s0
```

c. List information about the swap configuration.

In this example, swap is on the SVM root encapsulated mirror.

```

root@SourceSystem# swap -l
swapfile          dev      swaplo blocks    free
/dev/md/dsk/d20   85,20    16      133123696  133123696
root@SourceSystem# metastat -p d20
d20 -m d21 d22 1
d21 1 1  c0t1d0s1
d22 1 1  c1t0d0s1

```

Note - The `swap -l` command displays the swap space in 512-byte blocks. To convert the value to gigabytes, use the `bc` command. For example:

```

# bc
133123696*512/1024/1024/1024
63

```

d. Display information about the database binaries.

In this example, the binaries are on a UFS file system (`/u01`) on top of an SVM metadvice using two mirrored 150GB disks.

```

root@SourceSystem# df -h /u01
Filesystem          size  used  avail capacity  Mounted on
/dev/md/dsk/d30     148G  2.3G  144G      2%    /u01

root@SourceSystem# metastat -p d30
d30 -m d31 d32 1
d31 1 1  /dev/dsk/c2t600144F0E635D8C700005AC804B80019d0s0
d32 1 1  /dev/dsk/c2t600144F0E635D8C700005AC804D4001Ad0s0

root@SourceSystem# grep u01 /etc/mnttab
/dev/md/dsk/d30 /u01    ufs
    rw,intr,forcedirectio,largefiles,logging,xattr,onerror=panic,dev=154001e
    1528711266

root@SourceSystem# egrep "mount|u01" /etc/vfstab
#device      device      mount      FS    fsck    mount      mount
#to mount    to fsck     point     type  pass   at boot    options
/dev/md/dsk/d30 /dev/md/rdisk/d30 /u01      ufs    2      yes
    forcedirectio,logging

```

e. List the Solaris Volume Manager (SVM) metadvice state database (`metadb`) configuration.

The `metadb` is not migrated to the target system. Instead, `metadb` is recreated when preparing the target system (see [“Prepare the Target System” on page 32](#)).

```

root@SourceSystem# metadb
      flags          first blk      block count
a m p luo          16              8192      /dev/dsk/c0t1d0s4
a  p luo          8208              8192      /dev/dsk/c0t1d0s4
a  p luo          16400             8192      /dev/dsk/c0t1d0s4
a  p luo           16              8192      /dev/dsk/c1t0d0s4
a  p luo          8208              8192      /dev/dsk/c1t0d0s4
a  p luo          16400             8192      /dev/dsk/c1t0d0s4
    
```

f. Display information about the database files.

In this example, the Oracle database data files, and redo and archive log files are on a UFS file system on top of an SVM metadevice using two mirrored 200 GB disks.

```

root@SourceSystem# df -h /oradata
Filesystem          size  used  avail capacity  Mounted on
/dev/md/dsk/d40     197G   12G   182G     7%      /oradata
    
```

```

root@SourceSystem# metastat -p d40
d40 -m d41 d42 1
d41 1 1 /dev/dsk/c2t600144F0E635D8C700005AC804F4001Bd0s0
d42 1 1 /dev/dsk/c2t600144F0E635D8C700005AC8050C001Cd0s0
    
```

```

root@SourceSystem# grep oradata /etc/mnttab
/dev/md/dsk/d40 /oradata      ufs
rw,intr,forcedirectio,largefiles,logging,xattr,onerror=panic,dev=1540028
1528711266
    
```

```

root@SourceSystem# egrep "mount|oradata" /etc/vfstab
#device          device          mount          FS      fsck  mount  mount
#to mount        to fsck         point          type    pass  at boot options
/dev/md/dsk/d40 /dev/md/rdisk/d40 /oradata ufs 2 yes forcedirectio,logging
    
```

g. 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 32](#).

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 29](#).

This table lists the disk calculation for this example.

No. of LUNs	Total Storage Capacity	Used Storage	Data Mgt.	Contents
2	2 x 487 = 974 GB	4.9 GB (root)	UFS, SVM mirror	Root file system and swap space

No. of LUNs	Total Storage Capacity	Used Storage	Data Mgt.	Contents
		+ 64 GB (swap) = 69 GB		Note - The swap space is not captured in shared storage, and not migrated to the target system. Eventually it is recreated on the target system, so the space must be accounted for.
2	2 x 150 = 300 GB	2.3 GB	UFS on SVM metaset	Database binaries (/u01)
2	2 x 200 = 400 GB	12 GB	UFS on SVM metaset	Database files (/oradata)
1	n/a	n/a	SVM	metadb Note - Not captured in shared storage, and not migrated to the target system. Eventually recreated on the target system.
	1.6 TB (Minimum amount of space required on the target system to host the source.)	19.2 GB (Used to calculate the amount of required shared storage.)		Totals (rounded)

10. Obtain the encrypted password for the source system's root user.

This encrypted password is later used to access the root user on the target solaris10 branded zone when it is configured in [“Configure an Oracle Solaris 10 Branded Zone on the Target System” on page 40.](#)

```
root@SourceSystem# grep root /etc/shadow | nawk -F':' '{print $2}'
6Vq4iXheX6nm
```

11. Check the status of the Database and listener.

In this example, the source system is installed with Oracle 10.2.0.5.0.

Note – The second and third commands must be performed by the oracle user.

```
root@SourceSystem# cat /var/opt/oracle/oratab
md1:/u01/oracle/product/10.2.0/db_1:N

oracle1@SourceSystem $ /u01/oracle/product/10.2.0/db_1/bin/lsnrctl status
LSNRCTL for Solaris: Version 10.2.0.5.0 - Production on 25-JUL-2018 16:28:48
Copyright (c) 1991, 2010, Oracle. All rights reserved.
Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))
STATUS of the LISTENER
-----
Alias                                LISTENER
```

```

Version                TNSLSNR for Solaris: Version 10.2.0.5.0 - Production
Start Date             25-JUL-2018 10:51:40
Uptime                 0 days 5 hr. 37 min. 8 sec
Trace Level            off
Security               ON: Local OS Authentication
SNMP                   OFF
Listener Log File      /u01/oracle/product/10.2.0/db_1/network/log/listener.log
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=SourceSystem)(PORT=1521)))
Services Summary...
Service "md1" has 1 instance(s).
  Instance "md1", status READY, has 1 handler(s) for this service...
Service "md1XDB" has 1 instance(s).
  Instance "md1", status READY, has 1 handler(s) for this service...
Service "md1_XPT" has 1 instance(s).
  Instance "md1", status READY, has 1 handler(s) for this service...
The command completed successfully

```

```

oracle1@SourceSystem $ ptree|grep ora_|grep -v grep
11221 ora_pmon_md1
11223 ora_psp0_md1
11225 ora_mman_md1
11227 ora_dbw0_md1
11229 ora_dbw1_md1
11231 ora_dbw2_md1
...

```

12. Sample one table in the application schema (SOE).

As a sanity check, the same table is sampled again after the migration (in [“Start and Verify the Migrated Database” on page 54](#) to confirm data integrity.

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

```

COUNT(*)
-----
1429790

```

13. 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.

```

oracle@client-sys: $ sqlplus system/welcome1@//SourceSystem.us.example.com:1521/md1
SQL*Plus: Release 12.2.0.1.0 Production on Wed Jul 25 11:00:53 2018
Copyright (c) 1982, 2016, Oracle. All rights reserved.

Connected to:

```

Oracle Database 10g Enterprise Edition Release 10.2.0.5.0 - 64bit Production
With the Partitioning, OLAP, Data Mining and Real Application Testing options

```
SQL> select name,log_mode,open_mode from v$database ;
```

```
NAME          LOG_MODE      OPEN_MODE
-----
MD1           ARCHIVELOG    READ WRITE
```

```
SQL> select * from v$recover_file;
no rows selected
```

```
SQL> select name from v$datafile union select name from v$controlfile union select
member from v$logfile union select name from v$archived_log ;
```

```
NAME
-----
/oradata/md1/archivelog/MD1/archivelog/2018_05_25/o1_mf_1_2_fjhkyx3w_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_26/o1_mf_1_3_fomnzzjk_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_27/o1_mf_1_4_foqy2qjs_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_29/o1_mf_1_5_fovgmd0x_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_30/o1_mf_1_6_fozv6z9k_.arc
/oradata/md1/control01.ctl
/oradata/md1/control02.ctl
/oradata/md1/control03.ctl
/oradata/md1/redo01.log
/oradata/md1/redo02.log
/oradata/md1/redo03.log
/oradata/md1/sysaux01.dbf
/oradata/md1/system01.dbf
/oradata/md1/undotbs01.dbf
/oradata/md1/users01.dbf
```

```
16 rows selected.
```

▼ 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 (FLAR) 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 the FLAR.

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 source system FLAR.

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)	19.2 ÷ 2	9.6 GB
If you are installing an Oracle Solaris patchset during the migration to the target system, add 14GB to the space allotment. This is mandatory for source systems running Oracle Solaris 10 Update 8 or earlier. Otherwise the patchset is optional.	9.6 + 14	23.6 GB
Add 30% (To provide additional headroom for the creation of temporary files, uncompressing files, and so on.)	23.6 + 7	30.6 GB

For this example scenario, a minimum of **12.8 GB** of available space is required on the shared storage.

b. Identify available storage in the target system.

The command output shows that there is approximately 1.09 TB available storage capacity, which is sufficient space to contain the capture of source system for this example scenario.

```
root@TargetGlobal# 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 global zone, create and export a file system.

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

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

In this example, TargetGlobal:/ovas is an NFS share exported from the target system.

```
root@SourceSystem# mkdir /ovas
root@SourceSystem# mount -F nfs TargetGlobal:/ovas /ovas
```

4. (If needed) Download an Oracle Solaris patchset to the shared storage.

If the source system is running Oracle Solaris 10 Update 8 or earlier, you must download a patchset. It will be used with the `zoneadm install -P` command, which will automatically install the patchset during the migration to the target system.

If the source system is running Oracle Solaris 10 Update 9 (or later), you can choose to have a patchset automatically installed during the migration to the target system, or not.

The particular patchset that you are able to obtain depends on your Oracle Service agreement. If possible, download the Recommended OS Patchset for Solaris 10 SPARC. Alternatively, you can download the CPU OS Patchset 2018/04 Solaris 10 SPARC patchset.

To understand the support policies and to download a patchset, go to <https://blogs.oracle.com/solaris/solaris-10-extended-support-patches-patchsets-released>. Additional information is also provided in MOS knowledge article with Doc ID 1272947.1.

Once downloaded, unzip the downloaded zip file.

▼ Prepare the Target System

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

- **CPU Resources** – During the shift, you can assign any amount of CPU resources to the `solaris10` branded zone that are appropriate for the 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 workload on the target system, then the target system should provide at minimum the same available CPU and memory resources that the `solaris10` branded zone had 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 branded zone on the target system, for example, if the target system has faster CPUs, then the target system can provide fewer CPU resources to the branded zone. In some cases, using fewer CPU cores reduces software licensing costs.

- **Memory Resources** – By default, the target `solaris10` branded zone is allocated the same amount of memory that it had 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 virtual disks and sizes on the target must match the source system's disk configuration, as described in this procedure.
- **Network Resources** – A network data link created from the global zone to be used by the `solaris10` branded zone as its network interface. Redundancy of links can leverage aggregations available from the global zone.

1. List the target system processors.

In this example, the target system has 16 cores.

```
root@TargetGlobal# psrinfo -pv
The physical processor has 8 cores and 64 virtual processors (0-63)
  The core has 8 virtual processors (0-7)
  The core has 8 virtual processors (8-15)
  The core has 8 virtual processors (16-23)
  The core has 8 virtual processors (24-31)
  The core has 8 virtual processors (32-39)
  The core has 8 virtual processors (40-47)
  The core has 8 virtual processors (48-55)
  The core has 8 virtual processors (56-63)
    SPARC-S7 (chipid 0, clock 4267 MHz)
The physical processor has 8 cores and 64 virtual processors (64-127)
  The core has 8 virtual processors (64-71)
  The core has 8 virtual processors (72-79)
```

```

The core has 8 virtual processors (80-87)
The core has 8 virtual processors (88-95)
The core has 8 virtual processors (96-103)
The core has 8 virtual processors (104-111)
The core has 8 virtual processors (112-119)
The core has 8 virtual processors (120-127)
SPARC-S7 (chipid 1, clock 4267 MHz)

```

2. List the amount of memory.

```

root@TargetGlobal# prtconf | grep Mem
Memory size: 260096 Megabytes

```

3. List the amount of swap space.

```

root@TargetGlobal# swap -hl
swapfile          dev          swaplo        blocks        free
/dev/zvol/dsk/rpool/swap 288,1        8K           128G          128G

```

4. Display the target system's network configuration.

```

root@TargetGlobal# dladm show-aggr -x
LINK      PORT          SPEED DUPLEX  STATE  ADDRESS          PORTSTATE
aggr0     --            10000Mb full  up     0:10:e0:d5:22:c2 --
          net0          10000Mb full  up     0:10:e0:d5:22:c2 attached
          net2          10000Mb full  up     0:10:e0:d5:22:c4 attached

```

```

root@TargetGlobal# dladm show-aggr
LINK      MODE POLICY  ADDRPOLICY          LACPACTIVITY LACPTIMER
aggr0     trunk L3,L4  auto                active        short

```

```

root@TargetGlobal# ipadm
NAME      CLASS/TYPE STATE  UNDER  ADDR
aggr0     ip        ok     --      --
  aggr0/v4 static    ok     --      192.0.2.45/23
lo0       loopback  ok     --      --
  lo0/v4  static    ok     --      127.0.0.1/8
  lo0/v6  static    ok     --      ::1/128
net4      ip        ok     --      --
  net4/v4 static    ok     --      203.0.113.77/24

```

```

root@TargetGlobal# netstat -rn

```

```

Routing Table: IPv4
  Destination          Gateway          Flags Ref    Use  Interface
-----
default                192.0.2.1       UG     4    26860570
192.0.2.0              192.0.2.45     U      8    25427791 aggr0

```

```
127.0.0.1          127.0.0.1          UH      4          576 lo0
203.0.113.0       203.0.113.77      U       3          2759126 net4
```

Routing Table: IPv6

Destination/Mask	Gateway	Flags	Ref	Use	If
::1	::1	UH	2	2	lo0

```
root@TargetGlobal# cat /etc/resolv.conf
#
# _AUTOGENERATED_FROM_SMF_V1_
#
# WARNING: THIS FILE GENERATED FROM SMF DATA.
# DO NOT EDIT THIS FILE. EDITS WILL BE LOST.
# See resolv.conf(4) for details.
search us.example.com examplecorp.com example.com
nameserver 198.51.100.198
nameserver 198.51.100.197
nameserver 198.51.100.132
```

5. Create a vNIC on top of the trunk aggregation.

```
root@TargetGlobal# dladm create-vnic -l aggr0 vnic0
```

```
root@TargetGlobal# dladm show-vnic
LINK          OVER          SPEED  MACADDRESS          MACADDRTYPE  IDS
vnic0         aggr0         10000  2:8:20:df:e0:73    random       VID:0
```

6. Display information about the target's global zone root file system.

```
root@TargetGlobal# zpool status rpool
pool: rpool
state: ONLINE
scan: none requested
config:
    NAME                STATE          READ WRITE CKSUM
    rpool                ONLINE        0    0    0
    mirror-0             ONLINE        0    0    0
        c0t5000CCA08040A058d0  ONLINE        0    0    0
        c0t5000CCA080416934d0  ONLINE        0    0    0
errors: No known data error
```

7. Use the `format` utility to partition devices that will be used for the UFS/SVM file systems in the `solaris10` branded zone.

This table shows the storage configuration that is created on the target system to support the incoming source system OS and workloads (determined in [“Review the Source System Configuration” on page 21](#)).

No. of LUNs	Total Storage Capacity	Data Mgt.	Contents
2	2 x 1 TB = 2 TB	ZFS mirror	Global zone rpool
2	2 x 487 GB = 974 GB	ZFS mirror	Oracle Solaris 10 branded zone zonepath
2	2 x 150 GB = 300 GB	UFS/SVM mirror	Database binaries (/u01)
2	2 x 200 GB = 400 GB	UFS/SVM mirror	Database files (/oradata)
1	2 x 1 GB = 2 GB	SVM	metadb in global zone Note - Recreated on the target system.
	3.7 TB		Totals (rounded)

8. Ensure that disk slices are configured to support all the storage items.

a. Use the `prtvtoc` command to check the slices for `metadb`.

In this example, two local disk slices (1 GB each) are used for the `metadb` disks:

```
c0t5000CCA08040B3BCd0s4
```

```
c0t5000CCA08040AB24d0s4.
```

```
root@TargetGlobal# prtvtoc /dev/dsk/c0t5000CCA08040B3BCd0s2|tail -3
* Partition Tag Flags Sector Count Sector Mount Directory
  2 5 01 0 2344108410 2344108409
  4 0 00 48195 2120580 2168774
```

```
root@TargetGlobal# prtvtoc /dev/dsk/c0t5000CCA08040AB24d0s2|tail -3
* Partition Tag Flags Sector Count Sector Mount Directory
  2 5 01 0 2344108410 2344108409
  4 0 00 48195 2120580 2168774
```

b. Use the `prtvtoc` command to check the slices for the database binaries.

In this example, the database binaries require two 150 GB iSCSI devices:

```
c0t600144F09F2C0BFD00005B36D6910058d0s0
```

```
c0t600144F09F2C0BFD00005B36D6A70059d0s0
```

```
root@TargetGlobal# prtvtoc /dev/rdisk/c0t600144F09F2C0BFD00005B36D6910058d0s2|tail -3
* Partition Tag Flags Sector Count Sector Mount Directory
```

```

0      2      00          0      314386468 314386467
2      5      01          0      314386468 314386467

```

```

root@TargetGlobal# prtvtoc /dev/rdisk/c0t600144F09F2C0BFD00005B36D6A70059d0s2|tail -3
* Partition Tag  Flags   Sector    Count      Sector      Mount Directory
0      2      00          0      314386468 314386467
2      5      01          0      314386468 314386467

```

c. Use the prtvtoc command to check the slices for the database REDO and archive logs.

In this example, the database REDO and archive logs require two 200 GB iSCSI devices:

```
c0t600144F09F2C0BFD00005B36D6C3005Ad0s0
```

```
c0t600144F09F2C0BFD00005B36D6D5005Bd0s0
```

```

root@TargetGlobal# prtvtoc /dev/rdisk/c0t600144F09F2C0BFD00005B36D6C3005Ad0s2|tail -3
* Partition Tag  Flags   Sector    Count      Sector      Mount Directory
0      2      00          0      419289484 419289483
2      5      01          0      419289484 419289483

```

```

root@TargetGlobal# prtvtoc /dev/rdisk/c0t600144F09F2C0BFD00005B36D6D5005Bd0s2|tail -3
* Partition Tag  Flags   Sector    Count      Sector      Mount Directory
0      2      00          0      419289484 419289483
2      5      01          0      419289484 419289483

```

d. Use the format command to review the disks.

```

root@TargetGlobal# echo|format
Searching for disks...done

```

```

AVAILABLE DISK SELECTIONS:
0. c0t5000CCA08040A058d0 <HGST-H101812SFSUN1.2T-A990-1.09TB>
   /scsi_vhci/disk@g5000cca08040a058
   /dev/chassis/SYS/HDD0/disk
1. c0t5000CCA080416934d0 <HGST-H101812SFSUN1.2T-A990-1.09TB>
   /scsi_vhci/disk@g5000cca080416934
   /dev/chassis/SYS/HDD1/disk
2. c0t5000CCA08040B3BCd0 <HGST-H101812SFSUN1.2T-A990 cyl 48638 alt 2 hd 255
   sec 189>
   /scsi_vhci/disk@g5000cca08040b3bc
   /dev/chassis/SYS/HDD2/disk
3. c0t5000CCA08040AB24d0 <HGST-H101812SFSUN1.2T-A990-1.09TB>
   /scsi_vhci/disk@g5000cca08040ab24
   /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. c0t600144F09F2C0BFD00005B36D6A70059d0 <SUN-ZFS Storage 7355-1.0 cyl 4873
alt 2 hd 254 sec 254>
/scsi_vhci/ssd@g600144f09f2c0bfd00005b36d6a70059
6. c0t600144F09F2C0BFD00005B36D6C3005Ad0 <SUN-ZFS Storage 7355-1.0 cyl 6499
alt 2 hd 254 sec 254>
/scsi_vhci/ssd@g600144f09f2c0bfd00005b36d6c3005a
7. c0t600144F09F2C0BFD00005B36D6D5005Bd0 <SUN-ZFS Storage 7355-1.0 cyl 6499
alt 2 hd 254 sec 254>
/scsi_vhci/ssd@g600144f09f2c0bfd00005b36d6d5005b
8. c0t600144F09F2C0BFD00005B36D6910058d0 <SUN-ZFS Storage 7355-1.0 cyl 4873
alt 2 hd 254 sec 254>
/scsi_vhci/ssd@g600144f09f2c0bfd00005b36d6910058
9. c0t600144F09F2C0BFD00005B9195D80079d0 <SUN-ZFS Storage 7355-1.0-300.00GB>
/scsi_vhci/ssd@g600144f09f2c0bfd00005b9195d80079
10. c0t600144F09F2C0BFD00005B9195EF007Ad0 <SUN-ZFS Storage 7355-1.0-300.00GB>
/scsi_vhci/ssd@g600144f09f2c0bfd00005b9195ef007a
11. c0t600144F09F2C0BFD00005B9020260078d0 <SUN-ZFS Storage 7355-1.0 cyl 3998
alt 2 hd 8 sec 32>
/scsi_vhci/ssd@g600144f09f2c0bfd00005b9020260078
Specify disk (enter its number): Specify disk (enter its number):

```

9. Install the Solaris Volume Manager package and reboot the system.

The Solaris Volume Manager package is needed for the metadb. In Oracle Solaris 11.3 and later, the package is not installed by default.

```
root@TargetGlobal# pkg install storage/svm
```

```
root@TargetGlobal# reboot
```

10. Create a mirrored zpool for the solaris10 branded zone ZFS root file system.

Note - Only ZFS is supported for zone root file systems in Oracle Solaris 11.3.

```
root@TargetGlobal# zpool create -m /zones/TargetS10bz TargetS10bz mirror
c0t600144F09F2C0BFD00005B9195D80079d0 c0t600144F09F2C0BFD00005B9195EF007Ad0
```

```
root@TargetGlobal# zpool status TargetS10bz
pool: TargetS10bz
state: ONLINE
scan: resilvered 73K in 1s with 0 errors on Fri Sep 7 10:15:54 2018
```

config:

NAME	STATE	READ	WRITE	CKSUM
TargetS10bz	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0

```

c0t600144F09F2C0BFD00005B9195D80079d0 ONLINE      0      0      0
c0t600144F09F2C0BFD00005B9195EF007Ad0 ONLINE      0      0      0

```

errors: No known data errors

11. Configure the Solaris Volume manager (SVM) metadevices.

a. Create the SVM metadvice database using redundant local disks.

```
root@TargetGlobal# metadb -a -c 3 -f /dev/dsk/c0t5000CCA08040B3BCd0s4
```

```
root@TargetGlobal# metadb -a -c 3 -f /dev/dsk/c0t5000CCA08040AB24d0s4
```

```

root@TargetGlobal# metadb
      flags      first blk      block count
      a          u          16            8192      /dev/dsk/
c0t5000CCA08040B3BCd0s4
      a          u          8208           8192      /dev/dsk/
c0t5000CCA08040B3BCd0s4
      a          u          16400          8192      /dev/dsk/
c0t5000CCA08040B3BCd0s4
      a          u          16            8192      /dev/dsk/
c0t5000CCA08040AB24d0s4
      a          u          8208           8192      /dev/dsk/
c0t5000CCA08040AB24d0s4
      a          u          16400          8192      /dev/dsk/
c0t5000CCA08040AB24d0s4

```

b. Create a 150 GB mirrored SVM metadvice.

```
root@TargetGlobal# metainit d31 1 1 c0t600144F09F2C0BFD00005B36D6910058d0s0
```

```
root@TargetGlobal# metainit -f d32 1 1 c0t600144F09F2C0BFD00005B36D6A70059d0s0
```

```
root@TargetGlobal# metainit d30 -m d31 d32
```

```

root@TargetGlobal# metastat d30
d30: Mirror
  Submirror 0: d31
    State: Okay
  Submirror 1: d32
    State: Okay
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 314386468 blocks (149 GB)
d31: Submirror of d30
  State: Okay

```

```

Size: 314386468 blocks (149 GB)
Stripe 0:
  Device                               Start Block  Dbase
State Reloc Hot Spare
  /dev/dsk/c0t600144F09F2C0BFD00005B36D6910058d0s0      0      No
Okay Yes
d32: Submirror of d30
State: Okay
Size: 314386468 blocks (149 GB)
Stripe 0:
  Device                               Start Block  Dbase
State Reloc Hot Spare
  /dev/dsk/c0t600144F09F2C0BFD00005B36D6A70059d0s0      0      No
Okay Yes
Device Relocation Information:
Device                               Reloc  Device ID
/dev/dsk/c0t600144F09F2C0BFD00005B36D6910058d0  Yes
  id1,ssd@n600144f09f2c0bfd00005b36d6910058
/dev/dsk/c0t600144F09F2C0BFD00005B36D6A70059d0  Yes
  id1,ssd@n600144f09f2c0bfd00005b36d6a70059

```

c. Create a mirrored SVM metadvice for /oradata.

```

root@TargetGlobal# metainit d41 1 1 c0t600144F09F2C0BFD00005B36D6C3005Ad0s0

root@TargetGlobal# metainit -f d42 1 1 c0t600144F09F2C0BFD00005B36D6D5005Bd0s0

root@TargetGlobal# metainit d40 -m d41 d42

root@TargetGlobal# metastat d40
d40: Mirror
  Submirror 0: d41
    State: Okay
  Submirror 1: d42
    State: Okay
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 419289484 blocks (199 GB)
d41: Submirror of d40
  State: Okay
  Size: 419289484 blocks (199 GB)
  Stripe 0:
    Device                               Start Block  Dbase
State Reloc Hot Spare
  /dev/dsk/c0t600144F09F2C0BFD00005B36D6C3005Ad0s0      0      No
Okay Yes
d42: Submirror of d40

```

```

State: Okay
Size: 419289484 blocks (199 GB)
Stripe 0:
  Device                               Start Block  Dbase
State Reloc Hot Spare
  /dev/dsk/c0t600144F09F2C0BFD00005B36D6D5005Bd0s0      0      No
Okay Yes
Device Relocation Information:
Device                               Reloc  Device ID
/dev/dsk/c0t600144F09F2C0BFD00005B36D6C3005Ad0      Yes
id1,ssd@n600144f09f2c0bfd00005b36d6c3005a
/dev/dsk/c0t600144F09F2C0BFD00005B36D6D5005Bd0      Yes
id1,ssd@n600144f09f2c0bfd00005b36d6d5005b

```

12. On the target global zone, adjust the swap space.

Configure the capped memory and swap values so that they are similar to the source system configuration.

```

root@TargetGlobal# swap -d /dev/zvol/dsk/rpool/swap
root@TargetGlobal# zfs set volsize=80G rpool/swap
root@TargetGlobal# swap -a /dev/zvol/dsk/rpool/swap

```

13. Create the UFS file systems for Oracle binaries and data files

```

root@TargetGlobal# newfs /dev/md/rdisk/d30
root@TargetGlobal# newfs /dev/md/rdisk/d40

```

▼ Configure an Oracle Solaris 10 Branded Zone on the Target System

1. Log into the target global zone as superuser.

2. Identify the target's CPU resources.

Run the [psrinfo\(1M\)](#) command in the global zone and identify the cores that can be assigned to the solaris10 branded zone.

```

root@TargetGlobal# psrinfo -tL
socket: 0 (lgroups: 1, 0)
core: 0
  cpus: 0-7
core: 1
  cpus: 8-15

```

```

core: 2
  cpus: 16-23
core: 3
  cpus: 24-31
core: 4
  cpus: 32-39
core: 5
  cpus: 40-47
core: 6
  cpus: 48-55
core: 7
  cpus: 56-63
socket: 1 (lgroups: 2, 0)
core: 8
  cpus: 64-71
core: 9
  cpus: 72-79
core: 10
  cpus: 80-87
core: 11
  cpus: 88-95
core: 12
  cpus: 96-103
core: 13
  cpus: 104-111
core: 14
  cpus: 112-119
core: 15
  cpus: 120-127

```

3. Use one of the following methods to create a configuration text file.

When the zone is created, the configuration file supplied the zone configuration parameters.

■ Create a configuration file from scratch.

This method works well for simple configurations. This is the method used in our example.

Create a text file that includes one `zonecfg` subcommand per line. Refer to the [zonecfg \(1M\)](#) man page for details.

In this example, cores 2, 3, 4 and 5 will be allocated to the `solaris10` branded zone. This can be done by specifying 2-5 or 2,3,4,5 for the `add dedicated-cpu` subcommand.

```

root@TargetGlobal# cat > /ovas/zonecfg.txt <<EOF

create -b
set brand=solaris10

```

```
set zonepath=/zones/TargetS10bz
set autoboot=true
set autoshutdown=shutdown
set ip-type=exclusive
add net
set physical=vnic0
end
add dedicated-cpu
set cores=2-5
end
add capped-memory
set physical=64G
set swap=64G
end
add fs
set dir=/u01
set special=/dev/md/dsk/d30
set raw=/dev/md/rdisk/d30
set type=ufs
add options forcedirectio
end
add fs
set dir=/oradata
set special=/dev/md/dsk/d40
set raw=/dev/md/rdisk/d40
set type=ufs
add options forcedirectio
end

EOF
```

■ **Create a configuring file using the `zonep2vchk -c` command.**

This method produces a configuration file that contains the resource limits and network configuration based on the source host's physical resources and networking configuration.

Note - Depending on the source configuration, a configuration file created by this method might result in a file with dozens of configuration parameters.

a. **On the source system, run the `zonep2vchk` command and redirect the output to a shared location.**

```
root@SourceSystem# zonep2vchk -c > /ovas/zonecfg.txt
```

b. **Edit the file to make the configuration appropriate for the target environment.**

4. **On the target global zone, configure an exclusive-IP solaris10 branded zone.**

```
root@TargetGlobal# zonecfg -z TargetS10bz -f /ovas/zonecfg.txt
```

5. **(If needed) Configure a system ID profile for the solaris10 branded zone.**

If the system configuration (host name, IP parameters, and so on) must change on target, configure a sysidcfg profile for the solaris10 branded zone that provides the new parameters. The configuration file is used during the shift to the target system in [“Shift the Workloads to the Target” on page 47](#).

In this example, the required target network parameters are derived from information collected in [“Prepare the Target System” on page 32](#).

The encrypted root password that you specify in this file was captured in [“Review the Source System Configuration” on page 21](#).

```
root@TargetGlobal# vi /ovas/sysidcfg_s10brandzone_profile
-----
system_locale=C
terminal=xterm
network_interface=vnic0 {
    hostname=TargetS10bz
    ip_address=192.0.2.120
    netmask=255.255.254.0
    default_route=192.0.2.1
    protocol_ipv6=no
}
security_policy=NONE
name_service=DNS {
    domain_name=us.example.com
    name_server=198.51.100.198,198.51.100.197,198.51.100.132
    search=us.example.com,examplecorp.com,example.com
}
timezone=US/Pacific
timeserver=localhost
nfs4_domain=dynamic
root_password=jh8xQURgtDG/w
auto_reg=disable
```


Lifting and Shifting the Workloads

These topics describe how to migrate the source system OS and workloads to an Oracle Solaris 10 branded zone on the target system:

- [“Lift the Workloads to the Shared Storage” on page 45](#)
- [“Shift the Workloads to the Target” on page 47](#)

▼ Lift the Workloads to the Shared Storage

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

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

```
oracle1@SourceSystem $ sqlplus / as sysdba
SQL*Plus: Release 10.2.0.5.0 - Production on Wed Aug 1 15:49:52 2018
Copyright (c) 1982, 2010, Oracle. All Rights Reserved.
Connected to:
Oracle Database 10g Enterprise Edition Release 10.2.0.5.0 - 64bit Production
With the Partitioning, OLAP, Data Mining and Real Application Testing options
```

```
SQL> shutdown immediate
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> exit
```

```
oracle1@SourceSystem $ lsnrctl stop LISTENER
```

```
LSNRCTL for Solaris: Version 10.2.0.5.0 - Production on 01-AUG-2018 15:53:26
Copyright (c) 1991, 2010, Oracle. All rights reserved.
```

```
Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))
The command completed successfully
```

2. **Create a FLAR of the source system.**

The `flarcreate(1M)` command creates a flash archive (FLAR) of the source system, including the root file system and the mounted UFS file systems. In this example, this includes the Oracle database installation and data. The FLAR is created on the shared storage so that it is easily accessed by the target system.

Syntax:

```
flarcreate -L pax -S -c -n name flar_pathname
```

Where:

- `-L pax` – Create the flash archive using portable archive interchange (pax) archiver.
- `-S` – Skip the disk space check and do not write the archive size data to the archive.
- `-c` – Compress the archive using the `compress(1)` command.
- `-n name`– Used to specify a name that describes the archive. *name* is supplied as the value of the `content_name` keyword. See `flash_archive(4)`.
- `flar_pathname` – Is the full path name of the created FLAR file.

Example:

Note - In this example, the `time` command is used to provide command duration information.

```
root@SourceSystem# time flarcreate -L pax -S -c -n my_s10_source_system /ovas/  
SourceSystem.flar  
Full Flash  
Checking integrity...  
Integrity OK.  
Running precreation scripts...  
Precreation scripts done.  
Creating the archive...  
Archive creation complete.  
Running postcreation scripts...  
Postcreation scripts done.  
Running pre-exit scripts...  
Pre-exit scripts done.  
  
real 1:03:42.9  
user 32:26.8  
sys 6:26.9  
  
root@SourceSystem# ls -l /ovas/SourceSystem.flar  
-rw-r--r-- 1 nobody nobody 23384788425 Aug 1 18:19 /ovas/SourceSystem.flar
```

▼ Shift the Workloads to the Target

This procedure deploys the FLAR on the target system.

1. **Log into the target global zone as superuser.**
2. **Verify that the new zone is configured.**

The zone was created in [“Prepare the Target System” on page 32](#).

```
root@TargetGlobal# zoneadm list -vc
ID NAME          STATUS    PATH                                BRAND  IP
0 global         running   /                                  solaris shared
- TargetS10bz   configured /zones/TargetS10bz solaris10 excl
```

3. **Install the workloads in the new zone using the FLAR that was created on the source system.**

This step uses the `zoneadm(1M)` command with the FLAR to restore the Oracle database installation and data into corresponding UFS file systems that are configured in the `solaris10` branded zone.

Syntax

```
zoneadm -z zonename install -a flar_pathname
```

Where:

- `-z zonename` – Specifies the name of the zone.
- `-a flar_pathname` – Specifies the full path name of the FLAR file.

In addition to the basic syntax, these `zoneadm` options enable you to control how the branded zone is created:

- `-p` – Preserves the system configuration parameters such as the node name and IP parameters. The `-p` and `-c` options are mutually exclusive.
- `-c profile_pathname` – Enables you to change the system configuration parameters (node name and IP parameters) by specifying the path to a file that provides the new configuration parameters (created in [“Configure an Oracle Solaris 10 Branded Zone on the Target System” on page 40](#)). The `-p` and `-c` options are mutually exclusive.
- `-P patchset_pathname` – Required for systems running Oracle Solaris 10 Update 8 or earlier. This option is only available on target systems running Oracle Solaris 11.4 SRU 12. This option applies the Recommended or CPU OS patchsets to the branded zone during the installation.

Examples

- **Preserving the system configuration parameters such as the node name and IP parameters (-p).**

```
root@TargetGlobal# zoneadm -z TargetS10bz install -a /ovas/SourceSystem.flar -p
```

- **Changing the system configuration parameters such as the node name and IP parameters (-c *profile_pathname*).**

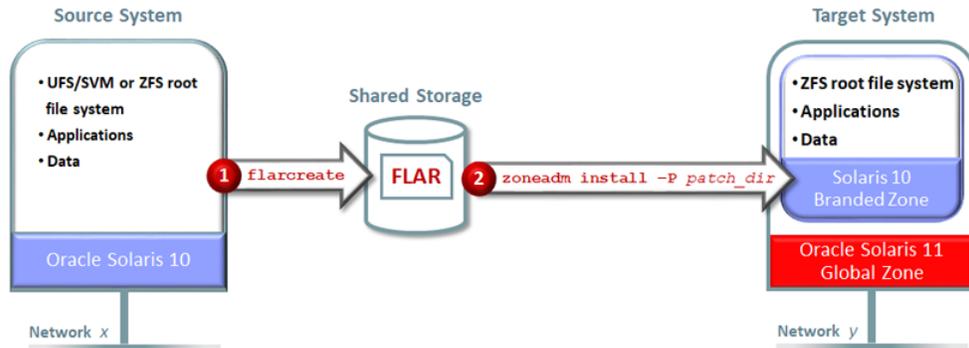
This is the method used for the example in this document.

```
root@TargetGlobal# zoneadm -z TargetS10bz install -a /ovas/SourceSystem.flar -u -
c /ovas/sysidcfg_s10brandzone_profile
Progress being logged to /var/log/zones/zoneadm.20180717T225458Z.TargetS10bz.install
  Installing: This may take several minutes...
Postprocessing: This may take a while...
  Postprocess: Updating the image to run within a zone
  Postprocess: Migrating data
    from: TargetS10bz/rpool/ROOT/zbe-0
    to: TargetS10bz/rpool/export
  Postprocess: A backup copy of /export is stored
    at /export.backup.20180717T230536Z.
It can be deleted after verifying it was migrated correctly.
```

```
      Result: Installation completed successfully.
Log saved in non-global zone as /zones/TargetS10bz/root/var/log/zones/
zoneadm.20180717T225458Z.TargetS10bz.install
```

- **Applying patchsets to the branded zone (-P *patchset_pathname*):**

The -P option is required for source systems running Oracle Solaris 10 Update 8 or earlier because Oracle Solaris 10 branded zones require an Oracle Solaris image that is equivalent to Oracle Solaris 10 Update 9 (or later). The patchsets applied with this option bring the zone image up to that level. If your source system is running Oracle Solaris 10 Update 9 or later, you can choose to use the -P option or not.



Recommended Patchset example:

```
root@TargetGlobal# zoneadm -z s10zone install -a /ovas/SourceSystem.flar -u -
P /ovas/10_Recommended
Progress being logged to /var/log/zones/zoneadm.20190730T185340Z.s10zone.install
```

Installing: This may take several minutes...

Zone image is running older release of Solaris, which cannot be run in solaris10 branded zone.

Patching: This may take a while ...

Patching: Applying prerequisite patches to the base ZBE ...

Patching: Applying 120900-04 (1 of 10) ... skipped

Patching: Applying 121133-02 (2 of 10) ... skipped

Patching: Applying 119254-93 (3 of 10) ... success

<output omitted for brevity>

Patching: Applying 152927-12 (419 of 421) ... skipped

Patching: Applying 152959-01 (420 of 421) ... success

Patching: Applying 153172-01 (421 of 421) ... skipped

Patching: Patch log files are saved in non-global zone directory:

Patching: /zones/s10zone/root/var/sadm/install_data/

Patching: /zones/s10zone/root/var/sadm/patch/

Postprocessing: This may take a while...

Postprocess: Updating the image to run within a zone

Postprocess: Migrating data

from: s10zone/rpool/ROOT/zbe

to: s10zone/rpool/export

Postprocess: A backup copy of /export is stored
at /export.backup.20190730T200153Z.
It can be deleted after verifying it was migrated correctly.

Result: Installation completed successfully.
Log saved in non-global zone as /zones/s10zone/root/var/log/zones/
zoneadm.20190730T185340.s10zone.install

Verifying the New Zone and Workloads

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

- [“Verify the Solaris 10 Branded Zone Configuration” on page 51](#)
- [“Start and Verify the Migrated Database” on page 54](#)
- [“Perform Assurance Tests and Release the Workloads into Production” on page 57](#)

▼ Verify the Solaris 10 Branded Zone Configuration

1. **Verify that the branded zone is installed.**

```
root@TargetGlobal# zoneadm list -vc
ID NAME          STATUS  PATH                                BRAND  IP
0 global         running /                                    solaris shared
- TargetS10bz    installed /zones/TargetS10bz                solaris10 excl
```

2. **In a separate terminal window, log into the zone's console.**

This console will display messages as the zone boots (next step).

```
root@TargetGlobal# zlogin -C TargetS10bz
```

Note - Once booted, you can use the `zlogin` command to log directly into the branded zone.

3. **(If needed) Shut down the source system.**

If the target branded zone uses the same network identity as the source system, shut down the source system to prevent conflicts on the network.

If the target branded zone is configured with a different network identity, you do not need to shut down the source system at this time.

4. **In the original terminal window, boot the branded zone.**

```
root@TargetGlobal# zoneadm -z TargetS10bz boot
```

5. Verify that the branded zone is running.

```
root@TargetGlobal# zoneadm list -vc
ID NAME          STATUS    PATH                                BRAND  IP
 0 global         running  /                                  solaris shared
 8 TargetS10bz   running  /zones/TargetS10bz                solaris10 excl
```

6. Log into the branded zone.

```
root@TargetGlobal# zlogin TargetS10bz
```

7. List the branded zone's system architecture.

```
rootTargetS10bz# uname -a
SunOS TargetS10bz 5.10 Generic_Virtual sun4v sparc sun4v
```

8. List the branded zone's processors

```
rootTargetS10bz# psrinfo -pv
The physical processor has 32 virtual processors (16-47)
SPARC-S7 (chipid 0, clock 4267 MHz)
```

9. List the branded zone's file systems.

```
rootTargetS10bz# df -h
Filesystem          Size  Used  Available Capacity  Mounted on
rpool/ROOT/zbe-0    293G  4.6G   275G      2%      /
rpool/ROOT/zbe-0/var 293G   14G   275G      5%     /var
/.SUNWnative/lib    892G  5.0G   887G      1%     /.SUNWnative/lib
/.SUNWnative/platform 892G  5.0G   887G      1%     /.SUNWnative/platform
/.SUNWnative/sbin   892G  5.0G   887G      1%     /.SUNWnative/sbin
/.SUNWnative/usr    892G  5.0G   887G      1%     /.SUNWnative/usr
/dev                 0K    0K     0K        0%     /dev
/oradata             197G  13G   182G      7%     /oradata
/u01                 148G  2.3G  144G      2%     /u01
proc                 0K    0K     0K        0%     /proc
ctfs                 0K    0K     0K        0%     /system/contract
mnttab               0K    0K     0K        0%     /etc/mnttab
swap                 64G  268M   64G      1%     /etc/svc/volatile
fd                   0K    0K     0K        0%     /dev/fd
swap                 64G  268M   64G      1%     /tmp
/etc/svc/volatile/   64G  268M   64G      1%     /var/run
rpool/export         293G   32K   275G      1%     /export
rpool/export/home    293G  5.2M   275G      1%     /export/home
```

```
rpool                293G   31K   275G   1%   /rpool
```

10. List the swap configuration.

```
rootTargetS10bz# swap -l
swapfile            dev  swaplo blocks   free
/dev/swap           0,0    16 134217728 133640208
```

11. Verify the branded zone's memory configuration.

```
rootTargetS10bz# prtconf |grep -i mem
prtconf: devinfo facility not available
Memory size: 65536 Megabytes
```

12. Verify the branded zone's network parameters.

```
rootTargetS10bz# ifconfig -a
lo0: flags=2001000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4,VIRTUAL> mtu 8232 index 1
    inet 127.0.0.1 netmask ff000000
vnic0: flags=100001000863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST,IPv4,PHYSRUNNING>
mtu 1500 index 2
    inet 192.0.2.120 netmask fffffe00 broadcast 192.0.3.255
    ether 2:8:20:df:e0:73
```

```
rootTargetS10bz# cat /etc/hosts
#
# Internet host table
#
127.0.0.1    localhost
::1        localhost
192.0.2.120 TargetS10bz    loghost
```

```
rootTargetS10bz# ls -l /etc/hostname*
-rw-r--r--  1 root    root          11 Jul 16 16:12 /etc/hostname.vnic0
```

```
rootTargetS10bz# cat /etc/hostname.vnic0
TargetS10bz
```

13. Log into the branded zone's database as the oracle user and check the database user environment.

```
rootTargetS10bz# su - oracle1
Oracle Corporation      SunOS 5.10      Generic Patch   January 2005
```

```
oracle1@TargetS10bz$ id
uid=100(oracle1) gid=10010(dba)
```

```
oracle1@TargetS10bz$ pwd
```

```
/export/home/oracle1  
  
oracle1@TargetS10bz$ ls  
db.env          oraInventory
```

▼ Start and Verify the Migrated Database

1. Login to the branded zone and switch to the oracle user.

```
root@TargetGlobal# zlogin TargetS10bz  
[Connected to zone 'TargetS10bz' pts/2]  
Last login: Wed Jul 16 05:05:11 on pts/5  
Oracle Corporation      SunOS 5.10      Generic Patch   January 2005  
  
root@TargetS10bz# su - oracle1  
Oracle Corporation      SunOS 5.10      Generic Patch   January 2005  
oracle1@TargetS10bz $
```

2. Adjust the listener configuration with a new host name.

Replace the old source host name with a new target guest domain name in the `listener.ora` parameter file.

```
oracle1@TargetS10bz$ cat $ORACLE_HOME/network/admin/listener.ora  
LISTENER =  
  (DESCRIPTION_LIST =  
    (DESCRIPTION =  
      (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC1))  
      (ADDRESS = (PROTOCOL = TCP)(HOST = TargetS10bz)(PORT = 1521))  
    )  
  )
```

3. Start the listener.

```
oracle1@TargetS10bz$ lsnrctl start  
  
LSNRCTL for Solaris: Version 10.2.0.5.0 - Production on 16-AUG-2018 06:05:08  
  
Copyright (c) 1991, 2010, Oracle. All rights reserved.  
  
Starting /u01/oracle/product/10.2.0/db_1/bin/tnslsnr: please wait...  
  
TNSLSNR for Solaris: Version 10.2.0.5.0 - Production  
Log messages written to /u01/oracle/product/10.2.0/db_1/network/log/listener.log  
Listening on: (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=TargetS10bz)(PORT=1521)))
```

```

Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))
STATUS of the LISTENER
-----
Alias                LISTENER
Version              TNSLSNR for Solaris: Version 10.2.0.5.0 - Production
Start Date           16-AUG-2018 06:05:08
Uptime               0 days 0 hr. 0 min. 0 sec
Trace Level          off
Security             ON: Local OS Authentication
SNMP                 OFF
Listener Log File    /u01/oracle/product/10.2.0/db_1/network/log/listener.log
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=TargetS10bz)(PORT=1521)))
The listener supports no services
The command completed successfully

```

4. Access and start the database.

```

oracle1@TargetS10bz$ sqlplus / as sysdba

SQL*Plus: Release 10.2.0.5.0 - Production on Fri Aug 16 05:06:45 2018

Copyright (c) 1982, 2010, Oracle. All Rights Reserved.

Connected to:
Oracle Database 10g Enterprise Edition Release 10.2.0.5.0 - 64bit Production
With the Partitioning, OLAP, Data Mining and Real Application Testing options

SQL> startup
ORACLE instance started.

Total System Global Area 1610612736 bytes
Fixed Size                2052448 bytes
Variable Size             385879712 bytes
Database Buffers          1207959552 bytes
Redo Buffers              14721024 bytes
Database mounted.
Database opened.

```

5. Check the status of the listener.

```

oracle1@TargetS10bz$ lsnrctl status

LSNRCTL for Solaris: Version 10.2.0.5.0 - Production on 16-AUG-2018 09:35:04

Copyright (c) 1991, 2010, Oracle. All rights reserved.

Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))

```

```

STATUS of the LISTENER
-----
Alias                LISTENER
Version              TNSLSNR for Solaris: Version 10.2.0.5.0 - Production
Start Date          16-AUG-2018 06:05:08
Uptime              0 days 3 hr. 30 min. 04 sec
Trace Level         off
Security            ON: Local OS Authentication
SNMP                OFF
Listener Log File   /u01/oracle/product/10.2.0/db_1/network/log/listener.log
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=TargetS10bz)(PORT=1521)))
Services Summary...
Service "md1" has 1 instance(s).
  Instance "md1", status READY, has 1 handler(s) for this service...
Service "md1XDB" has 1 instance(s).
  Instance "md1", status READY, has 1 handler(s) for this service...
Service "md1_XPT" has 1 instance(s).
  Instance "md1", status READY, has 1 handler(s) for this service...
The command completed successfully

oracle1@TargetS10bz$ ptree|grep ora_|grep -v grep
24088 ora_pmon_md1
24090 ora_psp0_md1
24092 ora_mman_md1
24094 ora_dbw0_md1
.....

```

6. Connect to the database from a client machine.

```

oracle@client-sys:$ sqlplus system/welcome1@//TargetS10bz.us.example.com:1521/md1

SQL> select name, log_mode,open_mode from v$database;

NAME          LOG_MODE      OPEN_MODE
-----
MD1           ARCHIVELOG    READ WRITE

SQL> select name from v$datafile union select name from v$controlfile union select
  member from v$logfile union select name from v$archived_log ;

NAME
-----
/oradata/md1/archivelog/MD1/archivelog/2018_05_25/o1_mf_1_2_fjhkyx3w_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_19/o1_mf_1_3_fo1fg0fd_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_20/o1_mf_1_4_fo5hgoso_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_22/o1_mf_1_5_fo8yh3d_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_23/o1_mf_1_6_fofdmzqy_.arc

```

```

/oradata/md1/archivelog/MD1/archivelog/2018_07_25/o1_mf_1_7_fojx3m2w_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_26/o1_mf_1_8_fonfm0cz_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_27/o1_mf_1_9_fopygmqj_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_28/o1_mf_1_10_fort6n3p_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_29/o1_mf_1_11_fow2j8gm_.arc
/oradata/md1/archivelog/MD1/archivelog/2018_07_30/o1_mf_1_12_fozv6qt6_.arc
/oradata/md1/control01.ctl
/oradata/md1/control02.ctl
/oradata/md1/control03.ctl
/oradata/md1/redo01.log
/oradata/md1/redo02.log
/oradata/md1/redo03.log
/oradata/md1/sysaux01.dbf
/oradata/md1/system01.dbf
/oradata/md1/undotbs01.dbf
/oradata/md1/users01.dbf
/u01/oracle/product/10.2.0/db_1/flash_recovery_area/MD1/archivelog/2018_05_24/o1
_mf_1_1_fjh62ty2_.arc

```

22 rows selected.

```
SQL> select * from v$recover_file;
```

no rows selected

7. **Check the application schema to verify that the expected data was migrated.**
The count should match the count from [“Review the Source System Configuration”](#) on page 21.

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

```

COUNT(*)
-----
1429790

```

▼ Perform Assurance Tests and Release the Workloads into Production

1. **Use the assurance tests for your environment to ensure the migrated zone is operable.**
2. **Shutdown the source system if it is still running.**
3. **Release the migrated workloads for use in your environment.**

Index

A

application schema, checking, 54
architecture, listing, 21

B

booting the zone, 51
branded zones, configuring, 40

C

capturing command output, 21
checking the source for migration suitability, 19
configuring
 branded zones on the target, 40
 SVM metadevices, 32
CPU resources, 32
creating
 flash archive (FLAR), 45
 zone configuration files, 40

D

database listener, starting, 54
databases on the target
 starting, 54
 verifying, 54
deploying the target FLAR, 47
description of the lift and shift scenario, 11
disks, listing, 21

F

flarcreate, 11
flash archive
 creating, 45
 deploying, 47

I

installing required patches, 21

L

lift and shift
 checking suitability, 19
 completing, 57
 overview, 11
 requirements, 15
 understanding the process, 9
lift and shift systems, preparing, 19
lifting the source to the shared storage, 45
listing
 source system architecture, 21
 source system disks, 21
 source system memory, 21
 source system metadevices, 21
 source system network parameters, 21
 source system processors, 21
 source system swap space, 21
logging into the zone console, 51

M

memory, listing, 21

metadevices, listing, 21
methods of migration, 9
migration methods based on the OS version, 9

N

network parameters, listing, 21

O

Oracle Solaris versions and migration methods, 9
ovas shared storage, 29
overview, lift and shift, 11

P

patches, installing, 21
preparing
 lift and shift systems, 19
 shared storage, 29
 target system, 32
processors, listing, 21
prompts used in examples, 16

R

requirements for the lift and shift, 15
reviewing
 source system, 21
 target system, 32

S

shared storage
 disk space calculation, 29
 preparing, 29
 requirements, 15
 transferring data to the target system, 47
 transferring the environment to, 45
shifting the source to the shared storage, 47
Solaris 10 branded zone configuration, verifying, 51

source system
 characteristics, 16
 prompt, 16
 requirements, 15
 reviewing, 21
starting
 database listener, 54
 migrated databases, 54
SVM metadevices, configuring on the target, 32
swap space, listing, 21
system ID profile, zone, 40

T

target system
 characteristics, 16
 disk names and capacities, 32
 preparing, 32
 prompt, 16
 requirements, 15
 verifying the zones, 51
testing the migration, 57

U

understanding the lift and shift process, 9

V

verifying
 migrated databases, 54
 Solaris 10 branded zone configuration, 51
 zones on the target system, 51
version dependencies, 9

Z

zone
 booting, 51
 configuring, 40
 system ID profile, 40

zone configuration file, creating, 40
zone console, logging into, 51
zoneadm, 47
zoneadm install, 11
zonep2vchk, 11, 19, 40

