Using Unified Archives for System Recovery and Cloning in Oracle Solaris 11.2

ORACLE®

Contents

Using This Documentation

- **Overview** – Describes the creation and deployment of Oracle Solaris Unified Archives to perform system recovery and cloning operations on Oracle Solaris systems
- **Audience** – System administrators and others responsible for system recovery and cloning of Oracle Solaris systems
- **Required knowledge** – Some experience working with Oracle Solaris

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at http://www.oracle.com/pls/topic/lookup?ctx=E36784.

Access to Oracle Support

Oracle customers 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.

Feedback

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

1

Oracle Solaris System Recovery and Cloning (Overview)

This book discusses the creation and deployment of Oracle Solaris Unified Archives to enable system recovery and cloning operations on Oracle Solaris systems. Oracle Solaris Unified Archives are system archives that can contain one or more instances of the operating system (OS). An OS instance may be a global zone, a non-global zone, or a kernel zone. These individual systems may be archived independently or bundled together. They may also be selectively archived, so that an archive may contain only one zone or a selection of zones.

This chapter covers the following topics:

About Oracle Solaris Unified Archives

Unified Archives are a new native archive type for Oracle Solaris. Unified Archives enable multiple system instances to be archived in a single unified file format. Unified Archives may contain one or more archived instances of Oracle Solaris from a single host. You can select individual installed zones to include during archive creation, and the host itself is optional.

You can deploy Unified Archives to recover a system that needs to be replaced due to failure, duplicate or clone a system configuration that you want installed on multiple machines, or migrate an existing system to new hardware or to a virtual system. You can deploy archives using any of the following methods:

- Oracle Solaris Automated Installer (AI)

- Oracle Solaris Zones utilities
- Unified Archive bootable media

Unified Archive Types

A Unified Archive type is a collection of archive attributes which serve a common use case. The two types of Unified Archive types are clone and recovery.

Both clone archives and recovery archives are created and managed by the `archiveadm` command. This command provides for selection of zones to be included or excluded from an archive, as well as allowing for ZFS dataset exclusion.

By default, an archive includes all datasets associated with the selected zones, unless they are specifically excluded during the archive creation process. For a non-global zone, the archive will contain the zone boot environment datasets and all delegated datasets assigned to it. For a global zone, the boot environment datasets are archived as well as all datasets not associated with a zone. Note that process includes entire top-down pools. Any number of datasets may be excluded from the archive. See "Unified Archives Creation" on page 10 for more information.

Clone Archives

A *clone archive* is based upon the system's active boot environment (BE) at the time the archive is created. Therefore, only one BE may be archived at a time. The clone archive type is intended for rapid creation and deployment of custom images. You can create an archive that includes both the global and non-global zones. The data for each archived system is stored independently, which enables each system to be independently deployable. To achieve independent deployment, clone relationships are not preserved within clone archives. Therefore, a cloned zone within a clone an archive will take up roughly the same amount of space as the original zone from which it is cloned. A clone archive is the default type of archive.

Also, a clone archive does not include any of the system configuration information from the OS instance, as well as any sensitive data such as SSH keys or passwords. Any system deployed from a clone archive will have all of the file systems and applications as the cloned system, but will not have the same host name, for instance. The system configuration information and sensitive data is only preserved in recovery archives.

Recovery Archives

A *recovery archive* is a full-system archive containing all boot environments from all included instances. Recovery archives consist of a single deployable system regardless of the OS

instances which comprise it. While a recovery archive of a global zone may contain non-global zones that are installed within it, the archive itself is treated as a single system. A recovery archive can therefore preserve zone clone relationships, unlike clone archives.

Some inactive boot environments in a recovery archive may not be bootable. The active boot environment is the only fully prepared boot environment. While some inactive BEs may be bootable, they are archived mainly for data recovery. Recovery archives contain a single deployable system. The system may include a global zone, a non-global zone, or a global zone with one or more non-global zones.

A recovery archive is intended to be used for recovery operations, as might be needed in case of a hardware failure. These archives may ideally be deployed as part of a more comprehensive disaster recovery solution. A recovery archive may also be used in a system migration from legacy hardware to new hardware or to virtual systems. Archives are portable across all supported platform models of a particular instruction set architecture (ISA).

Unified Archives Compared to Oracle Solaris 10 Flash Archives

Prior to the release of Oracle Solaris 11, the default archive solution was the Oracle Solaris Flash Archive. Introduced before the wide adoption of virtual systems, flash archives were designed to create and deploy OS instances of and to bare metal systems. Flash archives capture file system data from a running system along with system-related metadata. These archives are commonly used in pre-Oracle Solaris 11 environments to enable system recovery operations, as well as to achieve the creation and deployment of custom images.

With the introduction of boot environments (BEs), the Image Packaging System (IPS), and virtualization technologies systems such as Zones and LDOMs, a more robust archive solution is possible. Unified Archives employs these technologies and provides a much more flexible and extensible solution. It includes support for virtualized environments, such as zones, and cross-platform portability within the same hardware architecture.

Platform Portability and Unified Archives

Unified Archives are compatible between hardware platform models of the same instruction set architecture (ISA), as well as portable to virtual systems on the same ISA. For example, an archive created on a Sun SPARC Enterprise T5120 system is deployable to more modern hardware, such as the Oracle SPARC T5 line of systems, as well as Logical Domains or Solaris Kernel Zones running on newer hardware.

Archive Image Transforms

As an aspect of platform portability, Unified Archives enable image transforms between physical and virtual systems as well as global and non-global zones. Because these transforms are committed during deployment, the transform is determined by the deployed system type.

The most significant requirement for platform compatibility is the same instruction set architecture. Beyond this, any deployment target is supported for an archive, which enables many migration paths for legacy systems. For example, legacy hardware deployments may be archived and deployed as virtual system instances on new hardware. Similarly, non-global Oracle Solaris Zones may be archived and deployed as Oracle Solaris Kernel Zones.

There are some systems that can not be transformed, such as:

- Oracle Solaris 10 brand zones may not be transformed into any other system type. They may only be archived and deployed as an indentical Solaris 10 brand zone instance.
- Trusted Solaris or labeled zone archives can not be transformed. They can be deployed to a global zone contexts, such as a Logical Domain.

Unified Archives Creation

You can use the command-line utility `archiveadm` to create an archive. The `archiveadm` utility enables you to perform the following tasks:

- Create an archive to be used for system recovery or system cloning. See "Creating a Unified Archive" on page 16 for more information.
- Create bootable media from an archive to enable archive deployment. See "How to Deploy a System From a Unified Archive Using Bootable Media" on page 26 for more information.
- Retrieve information related to an existing archive, such as archive creation time, architecture, and system content. See "Viewing Unified Archive Information" on page 20 for more information.

For detailed procedures, see Chapter 2, "Working With Unified Archives".

Deployment Methods for Unified Archives

Once an archive has been created, you can store the archive as a file until it is needed. Deployment scenarios include system recovery and system migration, as well as system cloning. The three modes of archive deployment are:

- Non-interactive network deployment using the Automated Installer
- Zones deployment using the zones commands
- Interactive deployment using bootable media

Unified Archive Deployment With the Automated Installer

You can deploy archives by using the Oracle Solaris Automated Installer (AI). You have to modify the AI manifest to include the following information: the location of the archive and the name of the systems the archive should be deployed to. See "How to Deploy a System From a Unified Archive Using AI" on page 22 for more information.

Unified Archive Deployment With Zones

You use the Solaris Zones utilities `zonecfg` and `zoneadm` to configure and install new zones directly from an archive. When you create a new zone using an archive as a reference source, the new zone will mimic the original system's configuration. For specific instructions, see "How to Deploy a Zone From a Unified Archive" on page 23.

Zone creation and installation operations using archives are independent of each other. Although you can use the same archive to source both configuration and installation data, doing so is not a requirement. A zone configuration may already exist or may be created using any supported method.

Unified Archive Deployment With Bootable Media

You can deploy archives by bootable media using either ISO or USB images. You can use this form of deployment if the install client is not networked or if no AI services are available. These media images deploy the contents of the archive non-interactively when the client is booted.

You can use archives in any Oracle Solaris deployment scenario that uses bootable media, such as installing logical domains, Oracle Solaris Kernel Zones and Oracle Solaris instances within Oracle VM from media. For more information about bootable media installations, see "How to Deploy a System From a Unified Archive Using Bootable Media" on page 26.

Unified Archive Deployment to a Logical Domain

You can install logical domains from a Unified Archive either by using an Automated Installer service or by using bootable Unified Archive media. No special instructions are needed to deploy to a logical domain.

Unified Archive Deployment and AI Profiles

Deployment scenarios that support the use of Automated Installer (AI) profiles are supported with Unified Archives, including standard AI deployments, as well as Oracle Solaris Zones installation using the `zoneadm` command.

For archive types that retain system configuration information, this configuration will be overridden if a `sysconfig` profile is used. If a `sysconfig` profile is not used, the interactive system configuration screen will be invoked on the console for some archive deployments. For more information, see Chapter 11, "Configuring the Client System," in "Installing Oracle Solaris 11.2 Systems ".

Due to sysconfig's support for SMF configuration layering, some SMF configuration data may be retained in a recovery archive deployment even if a `sysconfig` profile is used to override it. Any information not overridden in a `sysconfig` profile will carry through into a deployment. For more information, see "Repository Layers" in "Managing System Services in Oracle Solaris 11.2 ".

Unified Archives and Trusted Solaris

Solaris Unified Archives support archive creation and deployment of Trusted Solaris and 'labeled'-brand Solaris Zones. This support is transparent in most respects save the following caveats:

- Archive image transforms are not allowed with Trusted Solaris or labeled zone archives. Therefore, Trusted Solaris global archives are deployable to global zone contexts such as logical domains, Oracle Solaris Kernel Zones, or to bare metal systems.
- Archives of labeled non-global zone archives are only deployable as 'labeled'-brand zones, running within the context of a Trusted Solaris global zone host. Note the requirement for 'labeled'-brand non-global zone deployments.
- Trusted Solaris must be configured in the host global zone. For more information regarding Trusted Solaris and 'labeled'-brand zones, refer to "Planning Your Labeled Zones in Trusted Extensions" in "Trusted Extensions Configuration and Administration ".

Best Practices Using Unified Archives

Thist section covers the following topics:

Improving Data Consistency With Multiple Pools

Archive components are consistent with one another, meaning that when an archive is deployed, all data present post-deployment should be from the same point in time from the origin system's perspective. Consistency is guaranteed on a per-pool basis. However, if an Oracle Solaris instance spans multiple pools, as when a root pool and a separate application or data pool are used. You should quiesce any active workloads prior to archive creation to ensure consistency between the pools once they are deployed.

Data Preparation

By default, all datasets other than swap and dump volumes are included when a Unified Archive is created. Datasets that are not needed for redeployment may be excluded using the exclude-dataset option (`-D`) to the `archiveadm create` command. Because directory and file exclusion lists are not provided, remove any temporary data or data that is otherwise not needed for redeployment prior to archive creation. When possible, place optional temporary data directories or any other data not needed for redeployment in their own datasets so that you can use the exclude-dataset option.

Naming Unified Archives

When creating Unified Archives, to use a file name that provides information about the archive. For example, `noscratch-recovery.uar` could tell you that the file is a recovery archive of a system that doesn't include the scratch dataset. Note that the name of the server that the archive was created from, the deployable systems in the archive, and the date the archive was created are stored in the metadata of the archive, which you can display using the `archiveadm info` command. You can see additional information by adding the `-v` option to this command. For examples showing the command output, see "Viewing Unified Archive Information" on page 20.

♦♦♦ C H A P T E R 2

2

Working With Unified Archives

This chapter discusses using Unified Archives for system cloning and system recovery deployments. It covers the following topics:

- "Using Rights to Control Access to Unified Archives" on page 15
- "Creating a Unified Archive" on page 16
- "Viewing Unified Archive Information" on page 20
- "Deploying Unified Archives" on page 21

For overview information about Unified Archives, see Chapter 1, "Oracle Solaris System Recovery and Cloning (Overview)".

Using Rights to Control Access to Unified Archives

The operating system includes five rights associated with creating and deploying Unified Archives:

- Install Manifest Management – All activities associated with creating and managing install manifests
- Install Profile Management – All activities associated with creating and managing install service profiles
- Unified Archive Administration – Can use the `archiveadm` command to create archives
- Zone Management – Can use the `zoneadm` command to to install or boot a zone and use the `zlogin` command to access the zone
- Zone Security – Can use the `zonecfg` command to create a new zone

▼ How to Configure a Role to Manage Unified Archives

For more information about using roles, see "Assigning Rights to Users" in "Securing Users and Processes in Oracle Solaris 11.2 ".

1. **Become an administrator.**

    For more information, see "Using Your Assigned Administrative Rights" in "Securing Users and Processes in Oracle Solaris 11.2 ".

2. **Create a new rights profile.**

    The following example adds two existing rights profiles to a new profile.

    ```
    # profiles -p new-profile
    profiles:new-profile> set desc="description"
    profiles:new-profile> add profiles="Unified Archive Administration"
    profiles:new-profile> add profiles="Zone Security"
    profiles:new-profile> commit; end; exit
    ```

3. **Create a role that uses the new rights profile.**

    In this step we also set a password for the new role.

    ```
    # roleadd -c "comment" -m -K profiles"new-profile" new-role
    # passwd new-role
    Password: xxxxxxxx
    Confirm password: xxxxxxxx
    ```

4. **Assign the new role to a user.**

    ```
    # usermod -R +new-role user
    ```

Creating a Unified Archive

You can use the `archiveadm create` command to create a Unified Archive. By default, a clone archive is created. Some useful options to the command include:

- `-r` to create a recovery archive
- `-z` to include specified zones
- `-Z` to exclude specified zones
- `-D` to exclude specified datasets

Once an archive is created, you can select from several methods to deploy a system using an archive. See "Deploying Unified Archives" on page 21 for more information. For a full list of all of the options see the `archiveadm`(1M) man page.

▼ How to Create a Clone Archive

1. **Become an administrator.**

    For more information, see "Using Your Assigned Administrative Rights" in "Securing Users and Processes in Oracle Solaris 11.2 ".

2. **Create the archive.**

    # **archiveadm create** *archive-name*

3. **(Optional) Verify the creation of the clone archive.**

    # **archiveadm info** *archive-name*

Example 2-1   Creating a Clone Archive That Includes All Zones

The following example shows how to create a clone archive with all zones into the clone archive named `production1.uar`.

```
# archiveadm create /var/tmp/production1.uar
Initializing Unified Archive creation resources...
Unified Archive initialized: /var/tmp/production1.uar
Logging to: /system/volatile/archive_log.22901
Executing dataset discovery...
Dataset discovery complete
Creating boot media for global zone(s)...
Media creation complete
Preparing archive system image...
Beginning archive stream creation...
Archive stream creation complete
Beginning final archive assembly...
Archive creation complete
# archiveadm info /var/tmp/production1.uar
Archive Information
          Creation Time:  2014-03-02T20:37:16Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
      Deployable Systems: global,example-net,example-dev
```

Example 2-2   Creating a Clone Archive That Includes Selected Zones

The following example shows how to create a clone archive named `zonearchive.uar` that includes the `example-net` and `example-dev` zones.

```
# archiveadm create -z example-net,example-dev /var/tmp/zonearchive.uar
Initializing Unified Archive creation resources...
 .
 .
Archive creation complete
```

```
# archiveadm info /var/tmp/zonearchive.uar
Archive Information
          Creation Time:  2014-03-02T17:04:11Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
      Deployable Systems: global,example-net,example-dev
```

The output from this command is the same as the text shown in Example 2-1 except for the name of the created archive and the log file name. Use the archiveadm info command to verify that the archive has the correct information.

Example 2-3   Creating a Clone Archive That Excludes ZFS Datasets

The following example demonstrates how to exclude a ZFS dataset when creating a clone archive.

```
# archiveadm create -D tank/scratch /var/tmp/noscratch.uar
 Initializing Unified Archive creation resources...
 .
 .
Archive creation complete
# archiveadm info /var/tmp/noscratch.uar
Archive Information
          Creation Time:  2014-03-02T17:04:11Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
      Deployable Systems: global,example-net,example-dev
```

The output from this command is the same as the text shown in Example 2-1 except for the name of the created archive and the log file name. Use the archiveadm info command to verify that the archive has the correct information.

▼ How to Create a Recovery Archive

1. **Become an administrator.**

   For more information, see "Using Your Assigned Administrative Rights" in "Securing Users and Processes in Oracle Solaris 11.2 ".

2. **Create the recovery archive.**

   ```
   # archiveadm create -r archive-name
   ```

3. **(Optional) Verify the creation of the recovery archive.**

   ```
   # archiveadm info -v archive-name
    ...
   ```

```
              Recovery Archive:  Yes
   ...
```

Example   2-4   Creating a Recovery Archive of the Global Zone

The following example shows how to create a recovery archive comprising only the global
zone.

```
# archiveadm create -r -z global /var/tmp/globalrecovery.uar
Initializing Unified Archive creation resources...
Unified Archive initialized: /var/tmp/globalrecovery.uar
Logging to: /system/volatile/archive_log.11234
Executing dataset discovery...
Dataset discovery complete
Creating boot media or global zone(s)...
Media creation complete
Preparing archive system image...
Beginning archive stream creation...
Archive stream creation complete
Beginning final archive assembly...
Archive creation complete
# archiveadm info globalrecovery.uar
Archive Information
          Creation Time:  2014-03-03T18:13:21Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
       Recovery Archive:  Yes
              ...

Deployable Systems:
          'global'
              ...
```

Example   2-5   Creating A Recovery Archive of a Zone

The following example shows how to create a recovery archive comprising a zone named
newzone. Note that only one zone is allowed for creation of a recovery archive.

```
# archiveadm create -r -z example-net /var/tmp/example-net-recovery.uar
Unified Archive initialized: /var/tmp/example-net-recovery.uar
 .
 .
Archive creation complete
# archiveadm info example-net-recovery.uar
Archive Information
          Creation Time:  2014-03-03T19:10:53Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
       Recovery Archive:  Yes
              ...

Deployable Systems:
```

```
        'example-net'
            ...
```

The output from this command is the same as the text shown in Example 2-4 except for the name of the created archive and the log file name. Use the `archiveadm info` command to verify that the archive has the correct information.

Example 2-6 Creating a Recovery Archive That Excludes a ZFS Dataset

As with clone archives, you can exclude datasets from an recovery archive. The following example shows how to create a recovery archive excluding the ZFS dataset `rpool/scratch`. The exclusion of a dataset is recursive, so all descendant datasets are also excluded.

```
# archiveadm create -r -D rpool/scratch /var/tmp/noscratch-recovery.uar
Unified Archive initialized: /var/tmp/noscratch-recovery.uar
 .
 .
Archive creation complete
# archiveadm info noscratch-recovery.uar
Archive Information
         Creation Time:  2014-03-03T19:46:10Z
           Source Host:  example
          Architecture:  i386
      Operating System:  Oracle Solaris 11.2 X86
      Recovery Archive:  Yes
            ...

Deployable Systems:
        'global'
            ...
```

The output from this command is the same as the text shown in Example 2-4 except for the name of the created archive and the log file name. Use the `archiveadm info` command to verify that the archive has the correct information.

Viewing Unified Archive Information

Use the `archiveadm info` command to examine Unified Archive file information. The examples in this section show both the abbreviated and verbose output.

EXAMPLE 2-7 Viewing Standard Information About an Archive

The following example shows the standard information displayed using the `archiveadm info` command.

```
% /usr/sbin/archiveadm info production1.uar
Archive Information
```

```
        Creation Time:  2013-10-02T20:37:16Z
          Source Host:  example
         Architecture:  i386
     Operating System:  Oracle Solaris 11.2 X86
    Deployable Systems: global,example-net,example-dev
```

EXAMPLE 2-8 Viewing All Information About an Archive

The following example shows the information displayed using the verbose option with the
`archiveadm info` command.

```
# archiveadm info -v production1.uar
Archive Information
          Creation Time:  2014-03-02T20:37:16Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86jQ
       Recovery Archive:  No
              Unique ID:  8700f573-34f5-c80b-a9c9-e7046523c6f4
        Archive Version:  1.0

Deployable Systems
        'global'
             OS Version:  5.11
              OS Branch:  5.11.0.0.0.26.2
              Active BE:  solaris
                  Brand:  solaris
            Size Needed:  6.1GB
              Unique ID:  27ff56cc-6c3e-ce90-8c7e-fb56fe79777f
      'example-net'
             OS Version:  5.11
              OS Branch:  5.11.0.0.0.26.2
              Active BE:  solaris
                  Brand:  solaris
            Size Needed:  684MB
              Unique ID:  6700df80-3a64-4aec-c500-907cdd2d648c
      'example-dev'
             OS Version:  5.11
              OS Branch:  5.11.0.0.0.26.2
              Active BE:  solaris
                  Brand:  solaris
            Size Needed:  684MB
              Unique ID:  a092fe9f-8319-c659-8e2c-803132af7ff7
```

Deploying Unified Archives

Once a Unified Archive has been created, you can deploy systems using either the Automated
Installer (AI) or bootable media. You can deploy a zone using the zones commands. For more
information, see:

- "How to Deploy a System From a Unified Archive Using AI" on page 22

- "How to Deploy a System From a Unified Archive Using Bootable Media" on page 26
- "How to Deploy a Zone From a Unified Archive" on page 23

Deploying a System From a Unified Archive Using AI

You must have a compatible AI service and AI manifest to deploy an archive using the Automated Installer (AI). A compatible service is one that is created from a version of the OS or an SRU, which is the same as or newer than the version from which the archive was created. For example, if the archived system is Oracle Solaris 11.2 FCS, any Oracle Solaris 11.2 version AI service is compatible. Each archive does not require its own service. The specific archive deployed is determined by the manifest used for a specific deployment.

To specify the Unified Archive location and image to deploy, use the ARCHIVE software type in an AI manifest. In Oracle Solaris 11.2, use the AI configuration `all_services/webserver_files_dir` property. This property sets a file system location that is served by the Automated Installer's web services and identifies a convenient location for files that support AI deployments.

For more information about customizing an AI manifest see Appendix A, "Example XML Manifests for Archive Recovery" and "Customizing an XML AI Manifest File" in "Installing Oracle Solaris 11.2 Systems ".

▼ How to Deploy a System From a Unified Archive Using AI

1. **Become an administrator.**

   For more information, see "Using Your Assigned Administrative Rights" in "Securing Users and Processes in Oracle Solaris 11.2 ".

2. **(Optional) If needed, create a new install service on the AI server.**

   It is only necessary to create a new install service if you do not already have one configured to support the needed client architecture and OS.

   # **installadm create-service -n** *new-service* **-s** *source*

   Ensure that the `-s` value is set to the name of the ISO file or the FMRI location for the current operating system version.

3. **Configure an AI manifest for the client.**

Either edit the default manifest XML file or create a new manifest XML file referencing the archive changes. Appendix A, "Example XML Manifests for Archive Recovery" provides sample manifests. Refer to "Customizing an XML AI Manifest File" in "Installing Oracle Solaris 11.2 Systems " for editing and verification guidelines.

4. **Create a new manifest with the updated manifest XML file.**

   # **installadm create-manifest -n** *svcname* **-m** *manifest* **-f** *file*

5. **Confirm both that the service was created and that the new manifest was applied to the service.**

   a. **Verify the new service.**

      # **installadm list -n** *new-service*

   b. **Verify the new manifest.**

      # **installadm list -n** *new-service* **-m**

Deploying a Zone From a Unified Archive

You can create and deploy Oracle Solaris Zones directly from Unified Archives. You can create new zone configurations using an archive as a template with the zonecfg command. You can install zones using zoneadm.

Note that zone configuration and zone installation are independent tasks. You can create new zone configurations from an archive and install the zone from IPS or other means. You can also use existing zone configuration and install using an archive. You can therefore source zone configurations and zone installations from different archives, as well. Refer to "Creating and Using Oracle Solaris Zones " for more information about zone configuration and installation.

▼ How to Deploy a Zone From a Unified Archive

1. **Become an administrator.**

   For more information, see "Using Your Assigned Administrative Rights" in "Securing Users and Processes in Oracle Solaris 11.2 ".

2. **If you do not already have an existing zone configuration, create one from an existing archive.**

   a. **Clone the zone configuration using the archive.**

Create a new zone configuration called new-zone using the zonecfg command. Use the second -z option to identify the zone to be cloned.

```
# zonecfg -z new-zone create -a archive -z archived-zone
```

b.   (Optional) Verify that the zone configuration was created.

```
# zoneadm list -cv
```

3.   Install the zone.

You can choose to install the zone using the software in an archive or using the software from the IPS repository.

a.   Install the zone using the archive.

```
# zoneadm -z new-zone install -a archive -z archived-zone
```

b.   Install the zone using the IPS repository.

```
# zoneadm -z new-zone install
```

4.   Boot the zone.

```
# zoneadm -z new-zone boot
```

Example 2-9   Configuring and Deploying a Zone

This example creates a clone of example-net named test-net. The zonecfg command shown below does not require a second -z option to define the zone to cloned since the example-net-recovery.uar archive only has one deployable system.

```
# archiveadm info /var/tmp/example-net-recovery.uar
Archive Information
          Creation Time:  2014-03-03T19:10:53Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
      Deployable Systems: example-net
# zonecfg -z test-net create -a /var/tmp/example-net-recovery.uar
# zoneadm list -cv
  ID NAME             STATUS      PATH                       BRAND     IP
   0 global           running     /                          solaris   shared
   - example-net      installed   /zones/example-net         solaris   excl
   - example-dev      installed   /zones/exmaple-dev         solaris   excl
   - test-net         configured  /zones/test-net            solaris   excl
# zoneadm -z test-net install -a /var/tmp/example-net-recovery.uar
Progress being logged to /var/log/zones/zoneadm.20131010T175728Z.test-net.install
       Image: Preparing at /zones/test-net/root.

 AI Manifest: /tmp/manifest.xml.caa4I2
```

```
      SC Profile: /usr/share/auto_install/sc_profiles/enable_sci.xml
        Zonename: test-net
Installation: Starting ...
   .
   .
Installation: Succeeded

        Note: Man pages can be obtained by installing pkg:/system/manual

 done.

        Done: Installation completed in 393.186 seconds.


  Next Steps: Boot the zone, then log into the zone console (zlogin -C)

               to complete the configuration process.

Log saved in non-global zone as /zones/test-net/root/var/log/zones/zoneadm.20131010T175728Z.test-
net.install
```

 # zoneadm -z test-dev boot

Example 2-10   Installing an Existing Zone Configuration From a Unified Archive

This example installs a zone called web-server using an existing zone configuration.

```
# archiveadm info /var/tmp/example-dev-recovery.uar
Archive Information
         Creation Time:  2014-03-03T19:10:53Z
           Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
     Deployable Systems: example-dev
# zoneadm list -cv
  ID NAME            STATUS      PATH                    BRAND     IP
   0 global          running     /                       solaris   shared
   - example-net     installed   /zones/example-net      solaris   excl
   - example-dev     installed   /zones/example-dev      solaris   excl
   - web-server      configured  /zones/web-server       solaris   excl
# zoneadm -z web-server install -a /data/archives/example-dev-recovery.uar
Progress being logged to  /var/log/zones/zoneadm.20140307T211655Z.web-server.install
    Installing: This may take several minutes...
 Install Log: /system/volatile/install.8799/install_log
 AI Manifest: /tmp/manifest.web-server.qCaakr.xml
    Zonename: web-server
Installation: Starting ...
   .
   .
 Updating non-global zone: Zone updated.
                   Result: Attach Succeeded.
       Done: Installation completed in 140.828 seconds.
 Next Steps: Boot the zone, then log into the zone console (zlogin -C)
             to complete the configuration process.
Log saved in non-global zone as
    /system/zones/web-server/root/var/log/zones/zoneadm.20140307T211655Z.web-server.install
```

```
# zoneadm -z web-server boot
```

Example   2-11   Configuring Zones From a Unified Archive With Multiple Deployable Systems

If multiple deployable systems are contained in a Unified Archive, use the -z option in zonecfg
to specify which deployable system to use. You can use only one deployable system.

```
# archiveadm info /var/tmp/zonearchive.uar
Archive Information
          Creation Time:  2014-03-03T17:04:11Z
            Source Host:  example
           Architecture:  i386
       Operating System:  Oracle Solaris 11.2 X86
      Deployable Systems: global,example-dev,example-net
# zonecfg -z test-dev create -a /var/tmp/zonearchive.uar -z example-dev
```

You can now install and boot the cloned zone.

Example   2-12   Installing Zones From a Unified Archive With Multiple Deployable Systems

If multiple deployable systems are contained in a Unified Archive, use the -z option in zoneadm
to specify the deployable system to install from.

```
# zoneadm -z test-dev install -a /var/tmp/zonearchive.uar -z example-dev
```

Deploying a System From a Unified Archive Using Bootable Media

When creating bootable media, the default file type that is created by the archiveadm create-
media command is USB. If the image is larger than 4GB, the USB type is forced.

▼ How to Deploy a System From a Unified Archive Using Bootable Media

1. **Become an administrator.**

   For more information, see "Using Your Assigned Administrative Rights" in "Securing Users
   and Processes in Oracle Solaris 11.2 ".

2. **Create a media image.**

   Create a media image from an existing archive.

   ```
   # archiveadm create-media archive-name
   ```

3. **Create bootable media.**

   ■ **ISO images – Burn the `.iso` file to a CD or DVD.**

   ■ **USB images – Use the `usbcopy` utility to copy the image to a USB flash drive.**

   Note - You can add this utility to your system by installing the `pkg:/install/distribution-constructor` package.

4. **Boot from the media.**

   Boot the system from the device that contains the boot image. A "hands-free" installation is performed. After the installation completes and the system is rebooted, you will be ask asked to provide configuration information for the system.

Example   2-13   Creating an ISO media image.

   To create an ISO image, you must include the `-f` option as shown below.

   ```
   # archiveadm create-media -f iso archive.uar
   ```

♦ ♦ ♦   **C H A P T E R   3**

3

Troubleshooting Cloning and Archiving

This section discusses common troubleshooting problems for cloning and archiving.

Common Error Messages Encountered in Cloning and Archiving

Issues may be encountered when creating and deploying Unified Archives. Most issues related to archive creation involve system configuration, available space for archive storage, and connectivity to the package repositories which are used to prepare the image for redeployment. Common deployment issues include target configuration and selection problems and failure to configure the deployed system.

In the case of archive creation or deployment errors, logging files are created. These files contain detailed data regarding the errors encountered. Note that with Solaris Kernel Zones, some log files may be within the specific zone.

Common Error Messages When Creating an Archive

`all installed zones excluded, at least one zone is required`

Solution: When using the `archiveadm create` command, add a zone to be included using the `-z` option.

`archives larger than 4GB can not use ISO format`

Solution: When using the `archiveadm create-media` command, use the USB format for archives larger than 4GB.

`estimated archive size larger than staging area /root`

Solution: When an archive is created, data is streamed to the directory where the archive is created. A capacity check is performed prior to archive creation and it may fail. Sometimes datasets have the `refreservation` property set which causes the capacity check to fail. In this case, it is possible for the pool to appear to have enough space for the archive, but the referred space reservation causes the failure. Note that ZFS volumes are created with a default `refreservation` which is equal to a little bit more than their size.

Reduce or remove the `refreservation` property value or run the command so that the archive is created in another dataset.

`excluded dataset includes active boot environment, dataset cannot be excluded`

Solution: When using the `archiveadm create` command, the dataset for the active boot environment cannot be excluded. Remove the dataset from the `-D` option list.

`FS type '`*fs-type*`' is not allowed via fs-allowed property in the zone configuration. Media cannot be created.`

Solution: When using the `archiveadm create-media` command, neither the `ufs` or `pcfs` file system type may be set as the `fs-allowed` property when creating a zone configuration.

Use the `zonecfg -z` *zone* `info` command to view the `fs-allowed` property for the zone.

`multiple zones not allowed for recovery archive`

Solution: When using the `archiveadm create` command, this error occurs when the `-z` option includes more than one zone. Instead, create a clone archive create a clone archive or create a recovery archive for each zone individually.

`no global zones found in archive`

When using the `archiveadm create-media` command, this error indicates that the archive does not contain a global zone. When using this command, the archive passed in needs to contain a global zone or a kernel zone archive.

`only a single host global zone is allowed`

Solution: When using the `archiveadm create-media` command, create media images for each global zone separately.

`output path is not writeable`

Solution: When using the `archiveadm create` command, check the permissions of the directory that you are trying to write the archive in.

```
required package system/boot/grub not installed
```

Solution: When using the `archiveadm create-media` command on an x86-based system, the system/grub/boot package must be installed. Install this package from a package repository using the following command: `pkg install system/grub/boot`.

```
'solaris-kz' branded zones must be running or excluded from archive creation
```

Solution: When using the `archiveadm create` command, either boot the branded zone or exclude the zone using the `-Z` option.

```
source is not a Solaris Automated installer ISO image
```

Solution: When using the `archiveadm create-media` command, select an ISO image to use with the `-s` option.

```
source ISO does not support archive operations
```

When using the `archiveadm create-media` command, the ISO image selected for media creation does not support unified archives. Create an ISO image on a system with Oracle Solaris 11.2 installed to create an ISO archive that supports unified archive operations.

```
Staging area is out of space
```

Solution: It is possible for the pool to fill or dataset quota to be met after the capacity check is performed. If other archives are being created or data is somehow otherwise being written to the pool, it may fill. As with all IO operations, capacity planning will help avoid this issue.

```
unable to find an AI image source
```

During archive or media creation, the AI image cannot be found and downloaded. Make sure the `install-image/solaris-auto-install` package is present on the available publishers.

```
Unable to revert packages. Please check publishers.
```

Solution: During archive creation, the archive package image is prepared for redeployment. Part of this process performs IPS package operations. It is likely that these operations will require data from the publishers be available. Therefore, the configured publishers must be available during archive creation.

```
unable to verify ISO image provided
```

During archive creation or media-creation, the AI ISO image cannot be mounted and verified. There may be a problem with the image, the mountpoint, or the LOFI device. Verify that each of these items is correct.

```
'zone' and 'exclude_zone' are mutually exclusive
```

Solution: When using the `archiveadm create` command, the same zone was listed with the -z option which lists zones to include when creating an archive and the -Z option which lists zones to exclude.

```
zone exclusion not allowed for recovery archive
```

Solution: When using the `archiveadm create` command, either create a clone archive with excluded zones or remove the -Z option to create a recovery archive.

```
zones must not be in incomplete or unavailable states for recovery archive creation
```

Solution: When using the `archiveadm create` command, you can exclude the incomplete or unavailable zones from the archive by using the -Z option.

Common Error Messages When Deploying an Archive

```
Failed Checkpoints: target-selection Checkpoint execution error: Unable to locate the
disk 'c1d0' on the system'
```

Solution: The archive which contains non-root data is being deployed without a target specification for non-root pools. Solaris Automated Installer will auto-select a boot device to install the root pool to, but non-root pool targets must be specified. Add targets for non-root pools to enable the deployment to proceed.

```
IPS publishers are required
```

Solution: Various types of archive deployments (for example, nonglobal-to-global zone transforms) require access to IPS publishers during deployment. Publishers are specified in an AI manifest and must be reachable during deployment from the client.

```
zone of this brand not in archive archive.uar: Unified archive not usable
```

Solution: When deploying an archive to a zone, the brand in the zone configuration is not compatible with the brand of the zones in the unified archive. Either select an archive with a matching brand or change the brand.

A
♦♦♦ **A P P E N D I X   A**

Example XML Manifests for Archive Recovery

Example XML Manifests

This chapter includes example manifests for archive deployment. For further information about installation manifests, refer to "Customizing an XML AI Manifest File" in "Installing Oracle Solaris 11.2 Systems ". Also refer to the aimanifest(1M) man page.

EXAMPLE   A-1   Manifest to Deploy a Global Zone From a Clone Archive

In this example, the referenced clone archive contains a single global zone that can be used to deploy a global zone.

```
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">
<auto_install>
  <ai_instance name="archive0">
    <target name="desired">
      <logical>
        <zpool name="rpool" is_root="true">
        </zpool>
      </logical>
    </target>
    <software type="ARCHIVE">
      <source>
        <file uri="http://example-ai.example.com/datapool/global.uar">
        </file>
      </source>
      <software_data action="install">
         <name>global</name>
      </software_data>
    </software>
  </ai_instance>
</auto_install>
```

EXAMPLE   A-2   Manifest to Deploy a Non-Global Zone From a Clone Archive

The following example shows a manifest deploying a non-global zone from a cloned archive.

```
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">
<auto_install>
  <ai_instance name="zone_default">
    <target>
```

```
          <logical>
            <zpool name="rpool">
            </zpool>
          </logical>
        </target>
        <software type="ARCHIVE">
          <source>
            <file uri="http:///example-ai.example.com/archives/zone-01.uar"/>
          </source>
          <software_data action="install">
            <name>zone-01</name>
          </software_data>
        </software>
      </ai_instance>
    </auto_install>
```

EXAMPLE   A-3   Manifest to Deploy a System Using a System Recovery Archive

The following example manifest demonstrates deploying a system using a system recovery
archive.

```
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">
<auto_install>
  <ai_instance name="archive0">
    <target name="desired">
      <logical>
        <zpool name="rpool" is_root="true">
        </zpool>
      </logical>
    </target>
    <software type="ARCHIVE">
      <source>
        <file uri="http://example-ai.example.com/recovery.uar">
        </file>
      </source>
      <software_data action="install">
         <name>*</name>
      </software_data>
    </software>
  </ai_instance>
</auto_install>
```

EXAMPLE   A-4   Manifest to Deploy an AI Client From a Non-global Zone Archive

To transform a non-global zone to a global zone requires additional data such as drivers,
platform software and so on. The IPS publishers that contain this information must be reachable
from the AI client during the deployment. In this example, an AI client called test2 will be
configured using information from the archive named /tmp/myarchive.uar and from packages
from http://pkg.oracle.com/solaris/release/. The location of the IPS repository is
defined in a second software declaration of type IPS.

```
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">
<auto_install>
```

```
 <ai_instance name="default">
  <target>
     <logical>
       <zpool name="rpool" is_root="true">
          <filesystem name="export" mountpoint="/export"/>
          <filesystem name="export/home"/>
       </zpool>
     </logical>
   </target>  ...
   <software type="ARCHIVE">
     <source>
       <file uri="/tmp/myarchive.uar">
       </file>
     </source>
     <software_data action="install">
        <name>test2</name>
     </software_data>
   </software>
  <software type="IPS">
    <source>
      <publisher name "solaris">
        <origin name="http://pkg.oracle.com/solaris/release/">
      </file>
    </source>
  </software>
 </ai_instance>
</auto_install>
```

EXAMPLE A-5   Manifest to Deploy a Global Zone and Configure a Zone From the Same Archive

In addition to deploying a zone, you can also configure an additional zone from an archive by
adding the following line to the manifest:

```
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">
<auto_install>
  <ai_instance name="zone1">
    <target>
      <logical>
        <zpool name="rpool" is_root="true">
        </zpool>
      </logical>
    </target>
    <software type="ARCHIVE">
      <source>
       <file uri="http://myserver.com/archives/mysystem-clone.uar">
       </file>
      </source>
      <software_data action="install">
        <name>zone1</name>
      </software_data>
    </software>
    <configuration type="zone" name="zone1" source="archive:myzone"
      archive="http://myserver.com/archives/mysystem-clone.uar"/>
  </ai_instance>
</auto_install>
```

This will configure `zone1` based on the `myzone` instance in the `mysystem-clone.uar` file. Be default, `zone1` will be installed from the IPS repository on first reboot, after the global zone has been installed. If you wish to use the archive to install the zone as well, use the following command:

```
# installadm create-manifest -n <service> -f <zone1-manifest-path> -m zone1 -c zonename="zone1"
```

Index