

**Using Unified Archives for System
Recovery and Cloning in Oracle® Solaris
11.3**

ORACLE®

Part No: E54846
October 2017

Using Unified Archives for System Recovery and Cloning in Oracle Solaris 11.3

Part No: E54846

Copyright © 2014, 2017, 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: E54846

Copyright © 2014, 2017, 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
1 Oracle Solaris System Recovery and Cloning (Overview)	9
About Oracle Solaris Unified Archives	9
Unified Archive Types	10
Clone Archives	10
Recovery Archives	11
Unified Archives Compared to Oracle Solaris 10 Flash Archives	11
Platform Portability and Unified Archives	12
Archive Image Transforms	12
Unified Archives Creation	13
Deployment Methods for Unified Archives	13
Unified Archive Deployment With the Automated Installer	13
Unified Archive Deployment With Zones	14
Unified Archive Deployment With Bootable Media	14
Unified Archive Deployment to a Logical Domain	14
Unified Archive Deployment and AI Profiles	14
Unified Archives and Trusted Extensions	15
Best Practices Using Unified Archives	15
Improving Data Consistency With Multiple Pools	16
Data Preparation	16
Naming Unified Archives	16
SVR4 Packages and Clone Archives	16
2 Working With Unified Archives	19
Using Rights Profiles with Unified Archives	19
▼ How to Configure a Role to Manage Unified Archives	20
Creating a Unified Archive	20

▼ How to Create a Clone Archive	21
▼ How to Create a Recovery Archive	23
Viewing Unified Archive Information	26
Deploying Unified Archives	28
▼ How to Deploy a System From a Unified Archive Using AI	28
▼ How to Deploy a Zone From a Unified Archive	30
▼ How to Deploy a System From a Unified Archive Using Bootable Media	33
3 Troubleshooting Cloning and Archiving	35
Common Error Messages Encountered in Cloning and Archiving	35
Common Error Messages When Creating an Archive	35
Common Error Messages When Deploying an Archive	38
Common Error Messages When Creating Media from an Archive	39
A Example XML Manifests for Archive Deployment	41
Example XML Manifests	41
Index	47

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** – Experience administering an Oracle Solaris system

Product Documentation Library

Documentation and resources for this product and related products are available at <http://www.oracle.com/pls/topic/lookup?ctx=E53394-01>.

Feedback

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

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” on page 9](#)
- [“Unified Archive Types” on page 10](#)
- [“Unified Archives Compared to Oracle Solaris 10 Flash Archives” on page 11](#)
- [“Platform Portability and Unified Archives” on page 12](#)
- [“Unified Archives Creation” on page 13](#)
- [“Deployment Methods for Unified Archives” on page 13](#)
- [“Unified Archive Deployment With Bootable Media” on page 14](#)
- [“Best Practices Using Unified Archives” on page 15](#)

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 system. You can select individual installed zones to include during archive creation, and the system 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 systems, or

migrate an existing system to new hardware or to a virtual machine. 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 13](#) 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.

During archive creation, some content is reverted to its pristine state using package revert tags. During clone archive creation, the `system:dev-init` and `system:clone` revert tags are used. The effect is that device configuration is cleared. Clone archives have instance-specific information like log file content and some configuration files reverted to their initial state.

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).

During archive creation, some content is reverted to its pristine state using package revert tags. During recovery archive creation, the `system:dev-init` revert tag is used, which clears device configuration information.

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 machines, flash archives were designed to create and deploy OS instances of and to bare metal chines. 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 machines 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. The `archiveadm` command only works on data that is managed by ZFS.

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 machines 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 machines 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 machine 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 identical Solaris 10 brand zone instance.
- Trusted Extensions or labeled zone archives can not be transformed. They can be deployed to a global zone context, 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 20](#) 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 33](#) 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 26](#) 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 28](#) 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 30](#).

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 from 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 install 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 33](#).

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, “Defining AI Client System Configuration Parameters” in *Installing Oracle Solaris 11.3 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 11.3*](#).

Unified Archives and Trusted Extensions

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

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

Best Practices Using Unified Archives

This section covers the following topics:

- [“Improving Data Consistency With Multiple Pools” on page 16](#)
- [“Data Preparation” on page 16](#)
- [“Naming Unified Archives” on page 16](#)
- [“SVR4 Packages and Clone Archives” on page 16](#)

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 (-D)` to the `archiveadm create` command. If you exclude one or more datasets, the archive will be smaller and will be created faster. 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, 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 system 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 26](#).

SVR4 Packages and Clone Archives

The IPS packaging system makes certain tasks available during archive creation that weren't possible with SVR4 packages. For example, the creation process for recovery archives cleans

up `/dev` and prepares the mappings of physical device names to instance numbers in `/etc/path_to_inst`. The creation process for clone archives performs the same operations and also cleans up other data that is specific to an Oracle Solaris instance, such as removing system configuration information. Unified archives use tagging attributes defined in IPS manifests to perform these tasks. These attributes also enable users to tag their own IPS packages for use when creating an archive.

To take advantage of IPS functionality during archive creation, convert SVR4 packages to IPS packages and install the packages onto a ZFS dataset. The `archiveadm` command only works on data that is managed by ZFS. For more information, see [“Converting SVR4 Packages To IPS Packages”](#) in *Packaging and Delivering Software With the Image Packaging System in Oracle Solaris 11.3*.

Working With Unified Archives

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

- [“Using Rights Profiles with Unified Archives” on page 19](#)
- [“Creating a Unified Archive” on page 20](#)
- [“Viewing Unified Archive Information” on page 26](#)
- [“Deploying Unified Archives” on page 28](#)

For overview information about Unified Archives, see [Chapter 1, “Oracle Solaris System Recovery and Cloning \(Overview\)”](#).

Using Rights Profiles with Unified Archives

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

- **Install Manifest Management** – Use when you are creating and managing install manifests
- **Install Profile Management** – Use when you are creating and managing install service profiles
- **Unified Archive Administration** – Use when you are using the `archiveadm` command to create archives
- **User Management** – Use when you are creating clone archives with the `archiveadm` command
- **Zone Management** – Use when you are installing or booting a zone with the `zoneadm` command or when accessing a zone with the `zlogin` command
- **Zone Security** – Use when you are creating a new zone with the `zonecfg` command

For more information, see [“Determining Which Rights to Assign to Administrators” in *Securing Users and Processes in Oracle Solaris 11.3*](#) and [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

▼ How to Configure a Role to Manage Unified Archives

The following procedure shows how to add rights profiles into one profile that you can use to manage unified archives.

1. Assume the root role.

In order to be able to configure a rights profile, you must assume the root role. For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

2. Create a new rights profile.

The following example adds 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="User Management"
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:

- `-D` to exclude specified datasets

- -e to exclude bootable media
- -r to create a recovery archive
- -z to include specified zones
- -Z to exclude specified zones

Once an archive is created, you can select from several methods to deploy a system using an archive. See [“Deploying Unified Archives” on page 28](#) 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 11.3*](#).

2. Create the archive.

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

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

```
# archiveadm info archive-name
```

Example 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: 2015-03-02T20:37:16Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 X86
    Deployable Systems: global,example-net,example-dev
```

Example 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: 2015-03-02T17:04:11Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 X86
    Deployable Systems: example-net,example-dev
```

The output from this command is the same as the text shown in [Example 1, “Creating a Clone Archive That Includes All Zones,” on page 21](#) 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 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: 2015-03-02T17:04:11Z
```

```

Source Host: example
Architecture: i386
Operating System: Oracle Solaris 11.3 X86
Deployable Systems: global,example-net,example-dev

```

The output from this command is the same as the text shown in [Example 1, “Creating a Clone Archive That Includes All Zones,” on page 21](#) 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 4 Creating a Clone Archive Without Bootable Media

By default, bootable media is created for each OS release in an archive. This media is embedded in the archive and used to boot and install Kernel Zones for each global zone in the archive. This process uses the `pkg://solaris/install-image/solaris-auto-install` package. If the package isn't available, or if you want to make smaller archives, create a clone archive without bootable media by using the `-e` option.

```

# archiveadm create -e /var/tmp/production2.uar
Initializing Unified Archive creation resources...
Unified Archive initialized: /var/tmp/production2.uar
Logging to: /system/volatile/archive_log.23912
Executing dataset discovery...
Dataset discovery complete
Preparing archive system image...
Beginning archive stream creation...
Archive stream creation complete
Beginning final archive assembly...
Archive creation complete

```

▼ 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 11.3*](#).

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 5 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 -v globalrecovery.uar
Archive Information
    Creation Time: 2015-03-03T18:13:21Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 X86
    Recovery Archive: Yes
    ...

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

Example 6 Creating A Recovery Archive of a Zone

The following example shows how to create a recovery archive comprising a zone named example-net. 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 -v example-net-recovery.uar
Archive Information
    Creation Time: 2015-03-03T19:10:53Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 X86
    Recovery Archive: Yes
    ...

Deployable Systems:
    'example-net'
    ...
```

The output from this command is the same as the text shown in [Example 5, “Creating a Recovery Archive of the Global Zone,”](#) on page 24 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 7 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 -v noscratch-recovery.uar
Archive Information
    Creation Time: 2015-03-03T19:46:10Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 X86
    Recovery Archive: Yes
    ...

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

The output from this command is the same as the text shown in [Example 5, “Creating a Recovery Archive of the Global Zone,”](#) on page 24 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 8 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: 2017-08-02T16:56:20Z
      Source Host:   example
      Architecture:  i386
      Operating System: Oracle Solaris 11.3 X86
      Deployable Systems: global
```

EXAMPLE 9 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: 2017-08-02T16:56:20Z
      Source Host:   example
      Architecture:  i386
      Operating System: Oracle Solaris 11.3 X86
      Recovery Archive: No
      Unique ID:     0d3333d8-42fa-42b5-9216-b442b96d9280
      Archive Version: 1.0

Deployable Systems
'global'
      OS Version: 0.5.11
      OS Branch:  0.175.3.23.0.3.0
      Active BE:  solaris
      Brand:      solaris
      Size Needed: 3.3GB
      Unique ID:  88e338aa-94e4-4754-8311-c16e0869e2f8
      AI Media:   0.175.3.22.0.3.0_ai_i386.iso
      Root-only:  Yes
```

EXAMPLE 10 Viewing Storage Configuration Information from the Origin System

The `archiveadm info` when used with the `-t` option (or the `--targets` option) shows you the storage configuration information from the system that the archive was created from.

```
# archiveadm info -t clone.uar
<target name="origin">
  <disk in_zpool="rpool" in_vdev="rpool-none" whole_disk="true">
    <disk_name name="/SYS/SASBP/HDD0" name_type="receptacle"/>
    <disk_prop dev_type="scsi" dev_vendor="HITACHI" dev_size="585937500secs"/>
    <disk_keyword key="boot_disk"/>
    <gpt_partition name="0" action="create" force="false" part_type="solaris">
      <size val="585920827secs" start_sector="256"/>
    </gpt_partition>
  </disk>
  <disk in_zpool="datapool1" in_vdev="datapool1-none" whole_disk="true">
    <disk_name name="/SYS/SASBP/HDD3" name_type="receptacle"/>
    <disk_prop dev_type="scsi" dev_vendor="HITACHI" dev_size="1172123568secs"/>
    <gpt_partition name="0" action="create" force="false" part_type="solaris">
      <size val="1172106895secs" start_sector="256"/>
    </gpt_partition>
  </disk>
  <logical noswap="false" nodump="false">
    <zpool name="datapool1" action="create" is_root="false" is_boot="false"
mountpoint="/datapool1">
      <vdev name="datapool1-none" redundancy="none"/>
    </zpool>
    <zpool name="rpool" action="create" is_root="true" is_boot="false" mountpoint="/
rpool">
      <vdev name="rpool-none" redundancy="none"/>
    </zpool>
  </logical>
</target>
```

You can take this information, then change it to match a new system in an AI manifest that can be used to deploy the system. In this example we changed the disk names.

```
<target name="origin">
  <disk in_zpool="rpool" in_vdev="rpool-none" whole_disk="true">
    <disk_name name="c1d0" name_type="ctd"/>
  </disk>
  <disk in_zpool="datapool1" in_vdev="datapool1-none" whole_disk="true">
    <disk_name name="c4d0" name_type="ctd"/>
  </disk>
  <logical noswap="false" nodump="false">
    <zpool name="datapool1" action="create" is_root="false" is_boot="false" mountpoint="
/datapool1">
      <vdev name="datapool1-none" redundancy="none"/>
    </zpool>
```

```
<zpool name="rpool" action="create" is_root="true" is_boot="false" mountpoint="/
rpool">
  <vdev name="rpool-none" redundancy="none"/>
</zpool>
</logical>
</target>
```

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 28](#)
- [“How to Deploy a System From a Unified Archive Using Bootable Media” on page 33](#)
- [“How to Deploy a Zone From a Unified Archive” on page 30](#)

▼ How to Deploy 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 11.3, any Oracle Solaris 11.3 or 11.3 SRU 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 11.3, 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 Deployment”](#) and [“Customizing an XML AI Manifest File” in *Installing Oracle Solaris 11.3 Systems*](#).

1. Become an administrator.

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

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. (Optional) If needed, associate the install client with a service.

In this example, the install client is associated with the new install service created in the previous step. The client is identified by its MAC address.

```
# installadm create-client -e ##.##.##.##.##.## -n new-service
```

4. Configure an AI manifest for the install 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 Deployment”](#) provides sample manifests. Refer to [“Customizing an XML AI Manifest File” in *Installing Oracle Solaris 11.3 Systems*](#) for editing and verification guidelines.

5. Create a new manifest with the updated manifest XML file.

```
# installadm create-manifest -n new-service -m manifest -f file
```

6. 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
```

7. Boot the install client.

This command will boot the install client using IP addresses from a DHCP server, which may also be your AI server.

```
ok boot net:dhcp - install
```

▼ How to Deploy 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.

- 1. Become an administrator.**

For more information, see “Using Your Assigned Administrative Rights” in [Securing Users and Processes in Oracle Solaris 11.3](#).

- 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 11 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: 2015-03-03T19:10:53Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 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/example-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 12 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: 2015-03-03T19:10:53Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 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.20150307T211655Z.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.20150307T211655Z.web-
server.install

# zoneadm -z web-server boot
```


Example 13 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: 2015-03-03T17:04:11Z
    Source Host: example
    Architecture: i386
    Operating System: Oracle Solaris 11.3 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 14 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
```

▼ How to Deploy 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.

- 1. Become an administrator.**

For more information, see [“Using Your Assigned Administrative Rights” in *Securing Users and Processes in Oracle Solaris 11.3*](#).

- 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-creator` 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 asked to provide configuration information for the system. For example, to boot from a CD, use the following command:

```
ok boot cdrom - install
```

Example 15 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
```

Example 16 Creating a Media Image Including a Specific Global Zone

By default the `create-media` subcommand will create media that includes all of the global zones that are included in the archive. To select a specific global zone, create the archive with the `-g global-zone` option. This command will create media from a specified zone. The resulting media image includes bootable media that contains the archive. This example will create an archive from the `global` global zone.

```
# archiveadm create-media -g global archive.uar
```

◆◆◆ CHAPTER 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 staging size of ##gb is larger than the available free space in /dataset. Additional space needed ##bg

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 reservation 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 reservation which is equal to a little bit more than their size.

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

excluded dataset dataset includes active boot environment, dataset 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`.

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 at least 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.

the following 'solaris-kz' branded zones must be running or excluded from archive creation: `zone1`, `zone2`

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

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 one or more non-root pools are being deployed without a target specification for those non-root pools. The Oracle Solaris Automated Installer will auto-select a boot device to install the root pool into, but non-root pool targets must be specified. Create an AI manifest that includes targets for non-root pools to enable the deployment to proceed.

IPS publishers are required for nonglobal-to-global transform deployments. Please provide an IPS specification.

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

Common Error Messages When Creating Media from an Archive

unable to find a supported AI image version (consider using '--source' option)

Solution: Running the `create-media` command requires access to a compatible AI image. If the archive was created with `-e` option to skip the generation of bootable media, the AI image is not embedded in the archive, so the command attempts to access the AI image package from a repository. If a compatible version of the AI image package is not available from the system's publishers, the command will fail. Use the `-s` or `--source` option to specify either a compatible ISO image or the repository URI where a compatible version of the AI image package can be found.

Example XML Manifests for Archive Deployment

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 11.3 Systems*. Also refer to the `aimanifest(1M)` man page. These sample manifests are provided:

- Example 17, “AI Manifest to Deploy a Global Zone From a Clone Archive,” on page 41
- Example 18, “AI Manifest to Deploy a Non-Global Zone From a Clone Archive,” on page 42
- Example 19, “AI Manifest to Deploy a System Using a System Recovery Archive,” on page 42
- Example 20, “AI Manifest to Deploy an AI Client From a Non-global Zone Archive,” on page 43
- Example 21, “AI Manifest to Deploy a Global Zone and Configure a Zone From the Same Archive,” on page 44
- Example 22, “AI Manifest to Deploy a Kernel Zone with Non-root Pools,” on page 45

Example XML Manifests

EXAMPLE 17 AI 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 18 AI 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 19 AI 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 20 AI 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. The IPS entry has to follow the ARCHIVE entry or the install will fail. The archive will not be deployed on the system if the package install happens first, so the ARCHIVE entry must be first.

```

<!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>
            </source>
        </software>
    </ai_instance>
</auto_install>

```

EXAMPLE 21 AI 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. By 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"

```

EXAMPLE 22 AI Manifest to Deploy a Kernel Zone with Non-root Pools

This manifest requires that devices with id's 0 and 1 exist in the zone's configuration.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE auto_install SYSTEM " file:///usr/share/install/ai.dtd.1">
<auto_install>
  <ai_instance name="default">
    <target name="origin">
      <disk in_zpool="rpool" in_vdev="rpool-none" whole_disk="true">
        <disk_name name="cld0" name_type="ctd"/>
      </disk>
      <disk in_zpool="tank" in_vdev="tank-none" whole_disk="true">
        <disk_name name="cld1" name_type="ctd"/>
      </disk>
      <logical noswap="false" nodump="false">
        <zpool name="rpool" action="create" is_root="true" mountpoint="/rpool">
          <vdev name="rpool-none" redundancy="none"/>
        </zpool>
        <zpool name="tank" action="create" is_root="false"
          mountpoint="/tank">
          <vdev name="tank-none" redundancy="none"/>
        </zpool>
      </logical>
    </target>
    <software type="ARCHIVE">
      <source>
        <file uri="file:///archives/clone.uar"/>
      </source>
      <software_data action="install">
        <name>*</name>
      </software_data>
    </software>
  </ai_instance>
</auto_install>
```


Index

A

AI

- deploying system, 28
- manifest examples, 41
- profiles
 - Unified Archives and, 14

archiveadm command

- creating a boot image, 33
- creating a clone archive, 21
- creating a recovery archive, 23
- introduction, 13
- viewing Unified Archive information, 26
- without including bootable media, 23

archiving

- overview, 9

B

best practices

- for Unified Archives, 15

boot image

- creating, 33

bootable media

- deploying system with, 33
- not including in a clone archive, 23

booting

- deploying a zone and, 30

C

clone archives

- creating, 21
- defined, 10

- excluding datasets, 22

- including all zones, 21

- SVR4 packages and, 16

- with selected zones, 22

- without bootable media, 23

cloning

- overview, 9

create subcommand

- archiveadm command example, 21, 23

create-manifest subcommand

- installadm command example, 28

create-media subcommand

- archiveadm command example, 33

create-service subcommand

- installadm command example, 28

creating

- boot image, 33

- clone archive, 21

- global zone recovery archive, 24

- install service, 28

- manifest, 28

- recovery archive, 23

- Unified Archives, 20

- zone recovery archive, 24

D

data preparation

- Unified Archives, 16

deploying

- logical domains, 14

- Unified Archives, 9

- using a Unified Archive with AI, 28

- using a Unified Archive with bootable media, 33
 - using a Unified Archive with zones commands, 30
- E**
- examples
 - AI manifest, 41
 - excluding datasets
 - from a clone archive, 22
 - from a recovery archive, 25
- F**
- f option
 - create-media subcommand, 34
 - Flash Archive
 - compared with Unified Archives, 11
- G**
- g option
 - create-media subcommand, 34
- I**
- improving data consistency
 - in Unified Archives, 16
 - info subcommand
 - archiveadm command example, 21, 23, 26
 - install service
 - creating, 28
 - installadm command
 - creating an install service, 28
 - installing
 - deploying a zone and, 30
 - ISO boot image
 - creating, 34
- L**
- logical domains
- deploying, 14
- M**
- manifest
 - creating, 28
 - examples, 41
- N**
- naming
 - Unified Archives, 16
- P**
- platform portability
 - Unified Archives and, 12
 - profiles
 - Unified Archives and, 14
- R**
- recovery archives
 - creating, 23
 - creating for a global zone, 24
 - defined, 10
 - for one zone, 24
- S**
- selecting
 - zone in media image, 34
 - set-service subcommand
 - installadm command example, 28
 - SVR4 packages
 - clone archives and, 16
- T**
- transforms

Unified Archives and, 12
Trusted Extensions
Unified Archives and, 15

deploying zone with, 30

U

Unified Archives
best practices, 15
compared to Flash Archive, 11
creating, 20
defined, 9
deploying using AI, 28
deploying using bootable media, 33
deploying using zones commands, 30
naming, 16
overview, 9
platform portability in, 12
Trusted Extensions and, 15
types, 10
viewing information about, 26

USB boot image
creating, 33

V

viewing
Unified Archive information, 26

Z

zoneadm command
deploying a system using, 30

zonecfg command
deploying a zone using, 30

zones
creating global zone recovery archive, 24
including all zones in a clone archive, 21
including selected zones in a clone archive, 22
specified in media image, 34

zones commands, 30, 30
See also zoneadm command
See also zonecfg command

