

Oracle® Solaris Cluster Data Service for Samba Guide

ORACLE®

Part No: E69338
February 2020

Part No: E69338

Copyright © 2000, 2020, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

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

Access to Oracle Support

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

Référence: E69338

Copyright © 2000, 2020, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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

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

Accès aux services de support Oracle

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

Contents

Using This Documentation	11
1 Installing and Configuring HA for Samba	13
HA for Samba Overview	13
Overview of Installing and Configuring HA for Samba	13
Planning the HA for Samba Installation and Configuration	14
Supported Configurations of Samba	15
Location of Samba Files	15
Requirements and Restrictions for Samba smb.conf File	15
Multiple Samba Instances That Require winbind	17
Dependencies Between HA for Samba Components	18
Required Parameters for the Samba smb.conf File	18
Installing and Configuring Samba Software	19
▼ How to Install and Verify Samba Software	20
▼ How to Prepare the Cluster for HA for Samba	22
▼ How to Configure Samba Software	23
Verifying the Configuration of Samba Software	25
▼ How to Verify the Configuration of Samba Software	26
Installing the HA for Samba Package	27
▼ How to Install the HA for Samba Package	27
Registering and Configuring HA for Samba	28
▼ How to Configure Resources for Samba and winbind	28
Verifying the HA for Samba Installation and Configuration	30
▼ How to Verify the HA for Samba Installation and Configuration	30
Understanding the HA for Samba Fault Monitor	31
Resource Properties	31
Probing Algorithm and Functionality	31
Debugging HA for Samba	32

- ▼ How to Turn On Debugging for HA for Samba 32

- A Deployment Example: Installing Samba Packaged with Oracle Solaris 11 Software in a Zone Cluster 35**
 - Target Cluster Configuration 35
 - Software Configuration 36
 - Deployment Example Assumptions 36
 - Configuring Samba Software 37
 - ▼ Example: Prepare the Cluster for Samba 38
 - ▼ Example: Configure the Logical Host and Storage Cluster Resources for Samba 38
 - ▼ Example: Create the Samba smb.conf Configuration File 39
 - ▼ Example: Create the Kerberos, PAM, and Name Service Switch Configuration for winbind 41
 - ▼ Example: Verify the Samba Configuration 42
 - ▼ Example: Enable Samba Software to Run in the Cluster 43
 - ▼ Example: Verify the HA for Samba Resource Group 45

- Index 47**

Tables

TABLE 1	Task Map for Installing and Configuring HA for Samba	14
TABLE 2	Samba Components and Their Dependencies	18

Examples

EXAMPLE 1	Running All Samba Instances and a winbind Instance Within the Same Failover Resource Group	17
EXAMPLE 2	Running Each Samba or winbind Instance Within Separate Failover Resource Groups	17

Using This Documentation

- **Overview** – Describes the Oracle Solaris Cluster HA for Samba data service.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

This document is not to be used as a planning or presales guide.

Product Documentation Library

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

Feedback

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

◆◆◆ CHAPTER 1

Installing and Configuring HA for Samba

This chapter explains how to install and configure the Oracle Solaris Cluster HA for Samba (HA for Samba) data service. This chapter contains the following sections.

- [“HA for Samba Overview” on page 13](#)
- [“Overview of Installing and Configuring HA for Samba” on page 13](#)
- [“Planning the HA for Samba Installation and Configuration” on page 14](#)
- [“Installing and Configuring Samba Software” on page 19](#)
- [“Verifying the Configuration of Samba Software” on page 25](#)
- [“Installing the HA for Samba Package” on page 27](#)
- [“Registering and Configuring HA for Samba” on page 28](#)
- [“Verifying the HA for Samba Installation and Configuration” on page 30](#)
- [“Understanding the HA for Samba Fault Monitor” on page 31](#)
- [“Debugging HA for Samba” on page 32](#)

HA for Samba Overview

The HA for Samba data service enables Oracle Solaris Cluster software to manage Samba by providing components to perform the orderly startup, shutdown, and fault monitoring of Samba.

HA for Samba can be used only with Samba software that is packaged with Oracle Solaris 11 software.

Overview of Installing and Configuring HA for Samba

The following table summarizes the tasks for installing and configuring HA for Samba and provides cross-references to detailed instructions for performing these tasks.

TABLE 1 Task Map for Installing and Configuring HA for Samba

Task	Instructions
Plan the installation.	“Planning the HA for Samba Installation and Configuration” on page 14
Install Samba software, configure the resource groups to be used by HA for Samba, and configure Samba software.	“Installing and Configuring Samba Software” on page 19
Verify Samba software installation and configuration.	“Verifying the Configuration of Samba Software” on page 25
Install the HA for Samba package.	“Installing the HA for Samba Package” on page 27
Register and configure HA for Samba resources.	“Registering and Configuring HA for Samba” on page 28
Verify HA for Samba installation and configuration.	“Verifying the HA for Samba Installation and Configuration” on page 30
Tune the HA for Samba fault monitor.	“Understanding the HA for Samba Fault Monitor” on page 31
Debug HA for Samba.	“Debugging HA for Samba” on page 32

Planning the HA for Samba Installation and Configuration

This section contains the information you need to plan your HA for Samba installation and configuration.

Throughout this section, references are made to the Samba instance and winbind instance. The HA for Samba data service consists of the components `smbd`, `nmbd`, and `winbindd`.

- The `smbd` and optional `nmbd` components are created within a single resource. This is referred to as the Samba instance or Samba resource.
- The `winbindd` component is created as a separate resource. This is referred to as the winbind instance or winbind resource.

For conceptual information about failover data services and scalable data services, see [Concepts for Oracle Solaris Cluster 4.4](#).

This section contains the following configuration requirements and restrictions for the HA for Samba data service.



Caution - Your data service configuration might not be supported if you do not observe these guidelines.

- [“Supported Configurations of Samba” on page 15](#)

- [“Location of Samba Files” on page 15](#)
- [“Requirements and Restrictions for Samba smb.conf File” on page 15](#)
- [“Multiple Samba Instances That Require winbind” on page 17](#)
- [“Dependencies Between HA for Samba Components” on page 18](#)
- [“Required Parameters for the Samba smb.conf File” on page 18](#)

Supported Configurations of Samba

HA for Samba supports Samba software in the following configurations.

- Primary Domain Controller (PDC).
- Backup Domain Controller (BDC) to a Samba PDC.
- NT4 domain member server, with or without winbind.
- Active Directory domain member server, with or without winbind.
- Standalone configuration.

HA for Samba is supported in the following Oracle Solaris Cluster configurations:

- The `smbd`, `winbindd`, and `nmbd` components can *only* be configured to run within a failover resource group.
- All components must run in the global zone or in a zone cluster. See [“Multiple Samba Instances That Require winbind” on page 17](#) for more information.

Location of Samba Files

The Samba files are where the Samba shares and `smb.conf` files are stored. The HA for Samba data service requires that these files be stored within a configuration directory that reflects the NetBIOS name for the Samba or winbind instance. The Samba files must be placed on shared storage, using either a cluster file system or a highly available local file system.

Requirements and Restrictions for Samba smb.conf File

The Samba `smb.conf` file is a configuration file that is used by the Samba and winbind instances.

The HA for Samba data service requires that such files be located at *configuration-directory/lib/smb.conf*. The following guidelines apply:

- Each Samba instance requires a unique *configuration-directory* that reflects the NetBIOS name of the Samba instance.
- A winbind instance can share a Samba instance *configuration-directory* and subsequent *smb.conf* file, together with the NetBIOS name of the Samba instance, if the Samba and winbind instances are both deployed in the same failover resource group.
- Each Samba instance *smb.conf* file must have a [scmondir] share. The HA for Samba fault monitor uses *smbclient* to access the directory specified within [scmondir] to verify that *smnd* is operating correctly. The following is an example of the [scmondir] share in an *smb.conf* file.

```
[scmondir]
comment = Monitor directory for Oracle Solaris Cluster
path = /tmp
browseable = No
```

The Samba instances run as failover services within separate failover resource groups on highly available local file systems with their own unique *configuration-directories*.

The following example uses the following component names:

- The NetBIOS name for the Samba instances are *smb1* and *smb2*.
- The NetBIOS name for the winbind instance is *winbind*.
- The Samba instance *smb1* has its *configuration directory* as */local/samba/config/smb1*.
- The Samba instance *smb2* has its *configuration directory* as */local/samba/config/smb2*.
- The winbind instance *winbind* has its *configuration directory* as */global/samba/config/winbind*.
- The Samba instance *smb1* has its *smb.conf* file located at */local/samba/config/smb1/lib/smb.conf*.
- The Samba instance *smb2* has its *smb.conf* file located at */local/samba/config/smb2/lib/smb.conf*.
- The winbind instance *winbind* has its *smb.conf* file located at */global/samba/config/winbind/lib/smb.conf*.

```
phys-schost# ls -l /global/samba/config
total 2
drwxrwx---  2 root    root      512 Jul 13 11:20 winbind
```

```
phys-schost# ls -l /local/samba/config
```



```
total 4
drwxrwx--- 2 root    root    512 Jul 13 11:25 smb1
drwxrwx--- 2 root    root    512 Jul 13 11:25 smb2
```

Multiple Samba Instances That Require winbind

The HA for Samba data service can support multiple Samba instances. However, only one winbind instance is supported per global zone or zone cluster.

If you intend to deploy multiple Samba instances that also require winbind, winbind needs to be a failover service. The following discussion will help you determine how to deploy single or multiple Samba instances with winbind. Disregard any reference to winbind if it is not required. Each example includes the following information:

- The zone cluster `samba-zc` is configured on two nodes of the global cluster, `node1` and `node2`.
- Benefits and drawbacks are listed within each example.

EXAMPLE 1 Running All Samba Instances and a winbind Instance Within the Same Failover Resource Group

This example creates a single failover resource group that contains all the Samba instances and a winbind instance in zone cluster `samba-zc` on global-cluster nodes `node1` and `node2`.

```
node1# clresourcegroup create -Z samba-zc RG1
```

Benefits

Only one non-global zone per node is required.

Drawbacks

All Samba/winbind instances do not have independent failover as they are all within the same failover resource group.

EXAMPLE 2 Running Each Samba or winbind Instance Within Separate Failover Resource Groups

Create multiple failover resource groups that will each contain one Samba or winbind instance in dedicated zone-cluster nodes across global-cluster nodes.

```
node1# clresourcegroup create -Z samba-zc1 RG1
node1# clresourcegroup create -Z samba-zc2 RG2
node1# clresourcegroup create -Z samba-zcN RGn
```

Benefits

All Samba or winbind instances have independent failover in separate failover resource groups.

All Samba or winbind instances are isolated within their own dedicated zone-cluster nodes.

Drawbacks

Each resource group requires a unique zone cluster.

Dependencies Between HA for Samba Components

If your Samba resource requires winbind, you must configure a start dependency on the winbind resource. After you create the Samba and winbind resources, you set this dependency in [“Registering and Configuring HA for Samba” on page 28](#).

[Table 2, “Samba Components and Their Dependencies,” on page 18](#) list the various dependencies.

TABLE 2 Samba Components and Their Dependencies

Component	Dependencies
Samba resource (smbd and nmbd)	winbind resource, if the Samba resource requires winbind services smbd logical host resource smbd HAStoragePlus resource
winbind resource (winbindd)	winbindd logical host resource winbindd HAStoragePlus resource

Note - Dependencies against the relevant component's logical host or HAStoragePlus resource are set for you when you register the Samba and winbind resources.

Required Parameters for the Samba smb.conf File

This section contains the following lists of required parameters for the Samba smb.conf file that is located within each configuration directory:

- [“Samba Parameters Required in smb.conf” on page 19](#)
- [“winbind Parameters Required in smb.conf” on page 19](#)

Refer to the `smb.conf(5)` man page for complete configuration information about Samba `smb.conf` file parameters.

Samba Parameters Required in `smb.conf`

- **bind interfaces only** – Must be set to yes.
- **interfaces** – Must be defined to the logical hostname.
- **lock dir** – Must include the *samba-configuration-directory* in its path.
- **netbios name** – Must be set to the NetBIOS name by which the Samba server is known.
- **pid directory** – Must include the *samba-configuration-directory* in its path.
- **security** – Specifies the security mode under which the Samba instance will run.
- **smb passwd file** – Must include the *samba-configuration-directory* in its path.

winbind Parameters Required in `smb.conf`

- **bind interfaces only** – Must be set to yes.
- **idmap**-related variables – Must be set to the same value as the Samba `smb.conf` entry.
- **interfaces** – Must be defined to the logical hostname.
- **lock dir** – Must include the *samba-configuration-directory* in its path.
- **netbios name** – Must be set to the NetBIOS name by which the winbind server is known.
- **password server** – Must be set to the same value as the Samba `smb.conf` entry.
- **pid directory** – Must include the *samba-configuration-directory* in its path.
- **template homedir** – Must be set to the same value as the Samba `smb.conf` entry.
- **template shell** – Must be set to the same value as the Samba `smb.conf` entry.
- **winbind enum groups** – Must be set to the same value as the Samba `smb.conf` entry.
- **winbind enum users** – Must be set to the same value as the Samba `smb.conf` entry.
- **winbind use default domain** – Must be set to yes.
- **workgroup** – Must be set to the same value as the Samba `smb.conf` entry.

Installing and Configuring Samba Software

This section provides procedures to install and configure Samba software.

Note - Oracle provides support for Samba that is packaged with Oracle Solaris 11 software, but does not offer support for Samba that has been downloaded and compiled from <https://www.samba.org>.

This section contains the following procedures:

- “How to Install and Verify Samba Software” on page 20
- “How to Prepare the Cluster for HA for Samba” on page 22
- “How to Configure Samba Software” on page 23

Note - The Oracle Solaris Cluster HA for Samba software can be configured to run in a zone cluster.

▼ How to Install and Verify Samba Software

Samba software is provided in the Oracle Solaris package `service/network/samba`. Perform the following procedure on each node to determine whether the same version of Samba software is installed on all nodes and to install Samba software if necessary.

- Before You Begin**
- Determine how many Samba instances to use. See “[Multiple Samba Instances That Require winbind](#)” on page 17.
 - Create a cluster file system or highly available local file system for the Samba files. For guidelines on locations for Samba files, see “[Location of Samba Files](#)” on page 15. For information about the file system for the Samba files, see [Chapter 9, “Updating Software Packages”](#) in *Updating Your Oracle Solaris Cluster 4.4 Environment*.

1. **Assume the root role.**
2. **Determine whether Samba software is already installed on any nodes in the cluster and, if so, which version is installed.**

Command output is similar to the following.

```
phys-schost-N# pkg info service/network/samba
Name: service/network/samba
Summary: samba - A Windows SMB/CIFS fileserver for UNIX
Category: System/File System
State: Installed
Publisher: solaris
Version: 3.6.12
```

```

Build Release: 5.11
Branch: 0.175.1.7.0.2.0
Packaging Date: Sat Apr 13 00:49:09 2013
Size: 185.70 MB
FMRI: pkg://solaris/service/network/samba@3.6.12,5.11-0.175.1.7.0.2.0:
20130413T004909Z

```

3. For any cluster node that is not already installed with the correct version of Samba software, install the Samba software package.

a. Ensure that the Samba package is available from the configured publisher and that the solaris publisher is valid.

```

# pkg list -a service/network/samba
# pkg publisher
PUBLISHER          TYPE    STATUS  P  LOCATION
solaris            origin  online  F  solaris-repository
ha-cluster        origin  online  F  ha-cluster-repository

```

For information about setting the solaris publisher, see [“Adding, Modifying, or Removing Package Publishers” in *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

Tip - Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the `-nv` options, run the command again without the `-n` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI to get additional information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,” in *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

b. Install the Samba software package.

```
phys-schost-N# pkg install service/network/samba
```

c. Verify that the package installed successfully.

```
$ pkg info service/network/samba
```

Installation is successful if output shows that State is Installed.

Next Steps Go to [“How to Prepare the Cluster for HA for Samba” on page 22](#).

▼ How to Prepare the Cluster for HA for Samba

Perform this procedure to prepare the cluster for the HA for Samba data service in a failover resource group.

Before You Begin Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.

1. **Assume a role that provides `solaris.cluster.modify` authorization.**
2. **Register the `SUNW.HAStoragePlus` and `SUNW.gds` resource types.**

```
phys-schost-1# clresource type register SUNW.HAStoragePlus
phys-schost-1# clresource type register SUNW.gds
```

3. **Create a failover resource group for Samba.**

```
phys-schost-1# clresourcegroup create -n nodelist samba-resource-group
```

4. **Create a resource for the Samba logical hostname.**

```
phys-schost-1# clreslogicalhostname create -g samba-resource-group \
-h samba-logical-hostname samba-logical-hostname-resource
```

5. **Create a resource for the Samba disk storage.**

- **If a ZFS highly available local file system is used, run the following command.**

```
phys-schost-1# clresource create -g samba-resource-group \
-t SUNW.HAStoragePlus \
-p Zpools=samba-zspool \
samba-hasp-resource
```

- **If a cluster file system or any other non-ZFS highly available local file system is used, run the following command.**

```
phys-schost-1# clresource create -g samba-resource-group \
-t SUNW.HAStoragePlus \
-p FilesystemMountPoints=samba-file-system-mountpoint \
samba-hasp-resource
```

6. **Bring online the failover resource group for Samba that now includes the `HAStoragePlus` and logical hostname resources.**

```
phys-schost-1# clresourcegroup online -M samba-resource-group
```

Next Steps Go to “[How to Configure Samba Software](#)” on page 23.

▼ How to Configure Samba Software

This section contains the steps to prepare Samba for use with the HA for Samba data service. Some steps require that you use Samba commands. Refer to the Docs and Books section under Learn Samba on the <https://www.samba.org> website for the relevant Samba man pages.

Perform this procedure on one node of the cluster, unless a specific step indicates otherwise.

Before You Begin Ensure that the same version of Samba software is installed on all nodes that you will configure with HA for Samba.

1. Create the fault monitor user.

- **If winbind is used, create the fault monitor user on the NT PDC or ADS server.**

- Use no home directory, no user profile, and no login script.
- Set the `Password never expire` parameter to `true`.
- Set the `User cannot change password` parameter to `true`.

- **If winbind is not used, perform the following commands on all nodes in the cluster.**

```
phys-schost-N# groupadd -g 1000 samba-fault-monitor-group
phys-schost-N# useradd -u 1000 -g 1000 -s /bin/false samba-fault-monitor-user
```

Note - A local Samba `fmuser` also requires a local password. The settings in the `smb.conf` file specify which password is used.

2. If winbind is used, activate the PAM configuration file on all nodes in the cluster.

```
phys-schost-N# cp -p /etc/pam.conf /etc/pam.conf.orig
phys-schost-N# cp /etc/pam.conf-winbind /etc/pam.conf
```

3. If winbind is used, configure the name service switch to resolve user and group from winbind on all nodes in the cluster.

```
phys-schost-N# svccfg -s name-service/switch setprop config/password = \"files winbind\"
phys-schost-N# svccfg -s name-service/switch setprop config/group = \"files winbind\"
phys-schost-N# svcadm refresh name-service/switch
```

4. **If winbind is used, disable the name service cache daemon on all nodes in the cluster.**

```
phys-schost-N# svcadm disable name-service-cache
```

5. **On one node of the cluster, create the Samba configuration directory.**

Perform this step for each Samba or winbind instance. Create the Samba configuration directory within the cluster file system or highly available local file system that is used for Samba files.

```
phys-schost-1# mkdir -p samba-configuration-directory
phys-schost-1# cd samba-configuration-directory
phys-schost-1# mkdir -p lib logs private shares var/locks
```

6. **On one node of the cluster, create the smb.conf file within the configuration directory that reflects the instance.**

Perform this step for each Samba or winbind instance.

Note - If security = share is required, you must include guest only = yes within [scomndir].

See [“Required Parameters for the Samba smb.conf File” on page 18](#) and the smb.conf(5) man page that is provided with the Samba software for an explanation of the required parameters.

7. **Add the NetBIOS name entry to the /etc/inet/hosts and /etc/inet/ipnodes files.**

Perform this step on each node that is used for Samba.

```
phys-schost-N# egrep -e "SMB1|ADS" /etc/inet/hosts /etc/inet/ipnodes
/etc/inet/hosts:192.168.1.132 SMB1#20
/etc/inet/hosts:192.168.1.9 ADS.EXAMPLE.COM#20
/etc/inet/ipnodes:192.168.1.132 SMB1#20
/etc/inet/ipnodes:192.168.1.9 ADS.EXAMPLE.COM#20
```

Note - The name resolve order parameter in the smb.conf file determines what naming service to use and in what order to resolve host names to IP addresses. Refer to the smb.conf [5] for more information.

The interfaces, netbios name, and password server parameters all require host name-to-IP address resolution.

8. **If Samba will operate as an Active Directory domain member server, create the Kerberos krb5.conf file.**

Perform this step on all nodes or zones that are used for Samba and that reflect the ADS realm. Refer to <https://www.samba.org> for complete information about installing and configuring Samba as a ADS domain member.

9. On one node, test the `smb.conf` file.

```
phys-schost-1# samba-bin-directory/testparm \
samba-configuration-directory/lib/smb.conf
```

10. If Samba is configured as an NT domain member, join the domain.

Perform this step from one node of the cluster.

```
phys-schost-1# samba-bin-directory/net -s samba-configuration-directory/lib/smb.conf \
RPC JOIN -U Administrator-on-the-PDC
```

11. If Samba is configured as a Windows 2003 domain member server with ADS, join the domain.

Perform this step from one node of the cluster.

```
phys-schost-1# samba-bin-directory/net -s samba-configuration-directory/lib/smb.conf \
ADS JOIN -U Administrator-on-the-ADS
```

12. If Samba is configured as a PDC or with `security = user`, add the fault monitor user.

Perform this step from one node of the cluster.

```
phys-schost-1# samba-bin-directory/smbpasswd \
-c samba-configuration-directory/lib/smb.conf \
-a samba-fault-monitor-user
```

13. If configured with `security = share`, verify that the `guest only` parameter is set to `yes`.

Inspect the `smb.conf` file to verify that `guest only = yes` is coded within the `[scondir]` section.

Next Steps Go to [“How to Verify the Configuration of Samba Software”](#) on page 26.

Verifying the Configuration of Samba Software

This section provides the procedure to verify the configuration of Samba software.

▼ How to Verify the Configuration of Samba Software

Perform this procedure from one node of the cluster.

1. Test the `smb.conf` file.

```
phys-schost-1# samba-bin-directory/testparm > samba-configuration-directory/Lib/smb.conf
```

2. If winbind is used, start and test winbind.

a. Start and test winbind.

```
phys-schost-1# samba-sbin-directory/winbindd \  
-s samba-configuration-directory/Lib/smb.conf  
phys-schost-1# getent passwd  
phys-schost-1# getent group
```

b. Test whether the fault monitor user can be resolved.

```
phys-schost-1# getent passwd samba-fault-monitor-user
```

This test must succeed. If you encounter problems, restart `winbindd` with debug information by using the `-d 3` option. Then retest and observe the `winbindd` log file, which is located at `samba-configuration-directory/logs/log.winbindd`.

Note - Winbind caching can affect the results from `getent passwd samba-fault-monitor-user`, which might not be up to date. See the `winbind(8)` man page for more information about winbind caching. See the `smb.conf(5)` man page for more information about winbind cache time.

3. Start and test Samba.

a. Start the Samba instance.

- If you use the `smbd` daemon, run the following command:

```
phys-schost-1# samba-sbin-directory/smbd \  
-s samba-configuration-directory/Lib/smb.conf -D
```

- If you use the `nmbd` daemon, run the following command:

```
phys-schost-1# samba-sbin-directory/nmbd \  
-s samba-configuration-directory/Lib/smb.conf -D
```

```
-s samba-configuration-directory/lib/smb.conf -D
```

b. Test whether the smbclient program can access Samba.

```
phys-schost-1# samba-bin-directory/smbclient -N -L NetBIOS-name
phys-schost-1# samba-bin-directory/smbclient '\\NetBIOS-name\scmondird' \
-U samba-fault-monitor-user -c 'pwd;exit'
```

This test must succeed. If you encounter problems, restart smbclient with debug information by using the `-d 3` option.

4. Stop the smbd, nmbd, and winbindd daemons.

Perform this step in the global zone only.

```
phys-schost-1# kill -TERM -z zonename 'smbd|nmbd|winbindd'
```

Next Steps Go to [“Installing the HA for Samba Package”](#) on page 27.

Installing the HA for Samba Package

If you did not install the HA for Samba package during your initial Oracle Solaris Cluster installation, perform this procedure to install the package.

▼ How to Install the HA for Samba Package

Perform this procedure on each cluster node where you want the HA for Samba software to run.

- 1. On the cluster node where you are installing the data service package, assume the root role.**
- 2. Ensure that the data service package is available from the configured publisher and that the solaris and ha-cluster publishers are valid.**

```
# pkg list -a ha-cluster/data-service/samba
# pkg publisher
PUBLISHER                TYPE      STATUS  P  LOCATION
solaris                   origin   online  F  solaris-repository
ha-cluster                 origin   online  F  ha-cluster-repository
```

For information about setting the solaris publisher, see [“Adding, Modifying, or Removing Package Publishers”](#) in *Updating Systems and Adding Software in Oracle Solaris 11.4*.

Tip - Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the `-nv` options, run the command again without the `-n` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,” in *Updating Systems and Adding Software in Oracle Solaris 11.4*](#).

3. Install the HA for Samba software package.

```
# pkg install ha-cluster/data-service/samba
```

4. Verify that the package installed successfully.

```
$ pkg info ha-cluster/data-service/samba
```

Installation is successful if output shows that State is Installed.

5. Perform any necessary updates to the Oracle Solaris Cluster software.

For instructions about updating your software, see [Chapter 9, “Updating Software Packages” in *Updating Your Oracle Solaris Cluster 4.4 Environment*](#).

Registering and Configuring HA for Samba

This section contains the procedure to configure HA for Samba resources.

▼ How to Configure Resources for Samba and winbind

Perform this procedure on one node of the cluster only.

- Before You Begin**
- Determine whether to deploy a single or multiple Samba instances, with or without winbind. See [“Multiple Samba Instances That Require winbind” on page 17](#).
 - Ensure that HA for Samba packages are installed. See [“How to Install the HA for Samba Package” on page 27](#).

1. **Assume a role that provides `solaris.cluster.modify` authorization.**
2. **If winbind is required, create, register, and enable a winbind resource.**
 - a. **Edit the `samba_config` file and follow the comments within that file.**
Ensure that `SERVICES="winbindd"` is specified.


```
phys-schost-1# cd /opt/SUNWscsmb/util
phys-schost-1# vi samba_config
```
 - b. **Register the resource.**

```
phys-schost-1# ./samba_register
```
 - c. **Enable the resource.**

```
phys-schost-1# clresource enable winbind-resource
```
3. **Configure an encrypted password for the Samba fault monitor.**
Perform this step on each node in the cluster.
 - a. **Set up the encrypted user password for the Samba fault monitor.**

```
phys-schost-N# touch /var/cluster/scsmb_key/opt/SUNWscsmb/.samba-lh_passwd
phys-schost-N# chmod 600 /var/cluster/scsmb_key/opt/SUNWscsmb/.samba-lh_passwd

phys-schost-N# dd if=/dev/urandom of=/var/cluster/scsmb_key bs=16 count=1
phys-schost-N# chmod 400 /var/cluster/scsmb_key

phys-schost-N# echo 'password' | /usr/sfw/bin/openssl enc -aes128 -e -pass \
file:/var/cluster/scsmb_key -out /opt/SUNWscsmb/.samba-lh_passwd
phys-schost-N# chmod 400 /opt/SUNWscsmb/.samba-lh_passwd
```
 - b. **Verify that the decrypted password matches the password for user `hasmb` within the Windows Active Directory domain.**

```
phys-schost-N# /usr/sfw/bin/openssl enc -aes128 -d \
-pass file:/var/cluster/scsmb_key -in /opt/SUNWscsmb/.samba-lh_passwd
```
4. **Create, register, and enable a Samba resource.**
 - a. **Edit the `samba_config` file and follow the comments within that file.**
Ensure that `SERVICES="smbd"` or `SERVICES="smbd,nmbd"` is specified.


```
phys-schost-1# cd /opt/SUNWscsmb/util
```

```
phys-schost-1# vi samba_config
```

b. Register the resource.

```
phys-schost-1# ./samba_register
```

c. Ensure that Samba is dependent on winbind.

```
phys-schost-1# clresource set -p Resource_dependencies=winbind-resource{local_node} \
samba-resource
```

d. Enable the Samba resource.

```
phys-schost-1# clresource enable samba-resource
```

Next Steps Go to [“Verifying the HA for Samba Installation and Configuration” on page 30.](#)

Verifying the HA for Samba Installation and Configuration

This section contains the procedure to verify that you installed and configured your data service correctly.

▼ How to Verify the HA for Samba Installation and Configuration

1. Assume a role that provides `solaris.cluster.modify` authorization.
2. Ensure that all Samba resources are online.

a. Check the status of Samba resources.

```
phys-schost-1# cluster status
```

b. Enable any Samba or winbind resource that is not online.

```
phys-schost-1# clresource enable samba-resource
```

3. Switch the Samba resource group to another cluster node.

```
phys-schost-1# clresourcegroup switch -n node samba-resource-group
```

Understanding the HA for Samba Fault Monitor

This section contains the following information.

- [“Resource Properties” on page 31](#)
- [“Probing Algorithm and Functionality” on page 31](#)

For conceptual information about fault monitors, see [Concepts for Oracle Solaris Cluster 4.4](#).

Resource Properties

The HA for Samba fault monitor uses the same resource properties as resource type `SUNW.gds`. See the [SUNW.gds\(7\)](#) man page for a complete list of resource properties.

Probing Algorithm and Functionality

The HA for Samba fault monitor is controlled by the extension properties that control the probing frequency. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Oracle Solaris Cluster installations. Therefore, you should tune the HA for Samba fault monitor *only* if you need to modify this preset behavior.

- Setting the interval between fault monitor probes (`Thorough_probe_interval`)
- Setting the time-out for fault monitor probes (`Probe_timeout`)
- Setting the number of times the fault monitor attempts to restart the resource (`Retry_count`)

The HA for Samba fault monitor checks the `smbd`, `nmbd`, and `winbindd` components in an infinite loop. During each cycle, the fault monitor checks the relevant component and reports either a failure or success.

- If the fault monitor is successful, it returns to its infinite loop and continues the next cycle of probing and sleeping.
- If the fault monitor reports a failure, a request is made to the cluster to restart the resource. If the fault monitor reports another failure, another request is made to the cluster to restart the resource. This behavior continues whenever the fault monitor reports a failure.

If successive restarts exceed the value of `Retry_count` in the `Thorough_probe_interval` property, a request is made to fail over the resource group onto a different node.

Operations of the winbind Probe

The winbind fault monitor periodically checks that the fault monitor user can be retrieved by using the `getent passwd samba-fault-monitor-user` command.

Note - The `winbindd` daemon resolves user and group information as a service to the name service switch. When running `winbindd`, the name service cache daemon must be turned off. To disable this daemon, see [Step 4](#) in “[How to Configure Samba Software](#)” on page 23.

Operations of the Samba Probe

The Samba probe checks the `nmbd` daemon by using the `nmblookup` program for each interface that is specified within the `smb.conf` file.

The Samba probe checks the `smbd` daemon by using the `smbclient` program, together with the `samba-fault-monitor-user`, to access the `scmond` share.

If the `smbclient` program cannot connect, there might be network or server issues that are causing the `smbclient` program to fail. These errors might be transient and correctable within a few seconds. Therefore, before a failure is called by the probe, the `smbclient` program is retried within 85% of the `Probe_timeout` property setting minus 15 seconds, which is approximately the timeout for the first `smbclient` failure. However, doing this is only realistic if the `Probe_timeout` setting is at least 30 seconds. If the `Probe_timeout` setting is less than 30 seconds, the `smbclient` program is tried only once.

Debugging HA for Samba

You can turn debugging on for all Samba or winbind instances or for only a particular Samba or winbind instance. A `config` file exists for this purpose. It is located in the `/opt/SUNWscsmb/instance/etc` directory, where *instance* refers to either `samba` or `winbind`.

If you need debugging enabled for HA for Samba across the entire cluster, perform the following procedure.

▼ How to Turn On Debugging for HA for Samba

Perform this procedure on each node in the cluster.

1. Display the `/etc/syslog.conf` file entries that contain the word "daemon".

```
phys-schost-N# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.notice;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                    operator
```

2. Open the `/etc/syslog.conf` file for editing and change `daemon.notice` to `daemon.debug`.

The following example output shows that `daemon.debug` is set.

```
phys-schost-N# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.debug;mail.crit      /var/adm/messages
*.alert;kern.err;daemon.err                    operator
```

3. Restart the `syslog` daemon.

```
phys-schost-N# svcadm restart system-log
```

4. On each cluster node that has Samba or winbind components that require debug output, enable debugging.

Edit the `/opt/SUNWscsmb/instance/config` file and set `DEBUG=` to either `DEBUG=ALL` or `DEBUG=resource`.

instance

The samba or winbind subdirectory.

ALL

Enables debugging for all resources of the Samba or winbind component.

resource

Enables debugging only for the specified resource.

The following example output shows a Samba configuration file with `DEBUG` set to `ALL`, to enable debugging on all Samba resources.

```
phys-schost-N# cat /opt/SUNWscsmb/samba/etc/config
#
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
#ident    "@(#)config 1.1    06/03/21 SMI"
#
# Usage:
#     DEBUG=<RESOURCE_NAME> or ALL
#
```

DEBUG=ALL

Next Steps To turn off debugging, reverse the steps in the previous procedure.

◆◆◆ A P P E N D I X A

Deployment Example: Installing Samba Packaged with Oracle Solaris 11 Software in a Zone Cluster

This appendix presents an example of how to install and configure the Samba application, shipped as part of the Oracle Solaris 11 IPS repository, and the HA for Samba data service in a zone cluster.

Note - To follow this example for configuration in a global cluster, perform these instructions from a global-cluster node instead of a zone-cluster node, and omit instructions for configuring a zone cluster.

This appendix contains the following information:

- [“Target Cluster Configuration” on page 35](#)
- [“Software Configuration” on page 36](#)
- [“Deployment Example Assumptions” on page 36](#)
- [“Configuring Samba Software” on page 37](#)

Target Cluster Configuration

This deployment example uses a two-node cluster that is configured with a two-node zone cluster. The nodes have the following names:

phys - node - 1

Physical node 1

phys - node - 2

Physical node 2

```
zc-node-1
  Zone cluster node 1

zc-node-2
  Zone cluster node 2
```

Software Configuration

This deployment example uses the following software products and versions:

- Oracle Solaris 11.2 software
- Oracle Solaris Cluster 4.2 software
- HA for Samba data service
- Samba shipped with Oracle Solaris 11.2 software
- Windows 2003 Enterprise Edition, with an Active Directory domain server configured

Deployment Example Assumptions

This deployment example is developed with the following assumptions:

- The two-node cluster is already installed and configured.
- The zone cluster `zc` is already installed and configured.
- The root user has the `/usr/cluster/bin` directory included in `$PATH` on all cluster nodes and zone cluster nodes.
- All cluster nodes have the `solaris` publisher for Oracle Solaris 11.2 software and the `ha-cluster` publisher for Oracle Solaris Cluster 4.2 software support repositories configured.
- The ZFS pool `samba_zp` is already created, using the mount point `/failover/samba`. The command used to create the ZFS pool is `zpool create -m /failover/samba samba_zp mirror devicelist mirror devicelist`.
- The zone cluster `zc` has the privileges to use the `zpool` and IP address within its zone cluster configuration. These privileges were set by using the following commands:

```
phys-node-1# clzonecluster configure zc
clzc:zc> add net
clzc:zc:net> set address=192.168.0.10
clzc:zc:net> end
clzc:zc> add dataset
```

```

clzc:zc:dataset> set name=samba_zp
clzc:zc:dataset> end
clzc:zc> commit
clzc:zc> exit

```

- The zone cluster nodes are configured as members of a Windows Active Directory domain, using winbind to resolve users and groups that are configured within the Active Directory domain.
- The Active Directory is configured to use the following values:

```

domain
    osc.example.com

```

```

domain NetBIOS name
    OSC

```

```

IP-Address
    192.168.0.20

```

- The following user and password are created within the Active Directory, to be used by the HA for Samba fault monitor:

```

Username
    hasmb

```

```

password
    xxxxxxxx

```

- No home directory, user profile, or login script are defined for the Active Directory.
- The Active Directory parameters Password never expire and User cannot change password are set to true.

Configuring Samba Software

This section contains the following tasks to install and configure Samba in a zone cluster:

- [“Example: Prepare the Cluster for Samba” on page 38](#)
- [“Example: Configure the Logical Host and Storage Cluster Resources for Samba” on page 38](#)
- [“Example: Create the Samba smb.conf Configuration File” on page 39](#)
- [“Example: Create the Kerberos, PAM, and Name Service Switch Configuration for winbind” on page 41](#)

- [“Example: Verify the Samba Configuration” on page 42](#)
- [“Example: Enable Samba Software to Run in the Cluster” on page 43](#)
- [“Example: Verify the HA for Samba Resource Group” on page 45](#)

▼ Example: Prepare the Cluster for Samba

1. **On all cluster nodes and zone-cluster nodes, add the logical hostname to the `/etc/inet/hosts` file.**

Ensure that you include the fully qualified domain name in an alias.

```
phys-node-N# grep samba-lh /etc/inet/hosts
192.168.0.10 samba-lh samba-lh.osc.example.com
```

2. **On all cluster nodes and zone cluster nodes, add the IP address for the Windows Active Directory server to the `/etc/inet/hosts` file.**

Ensure that you include the fully qualified domain name in an alias.

```
phys-node-N# grep windows-ad /etc/inet/hosts
192.168.0.20 windows-ad windows-ad.osc.example.com
```

3. **On all zone cluster nodes, install the Samba IPS package from the Oracle Solaris 11.2 repository.**

```
zc-node-N# pkg install --accept service/network/samba
```

4. **On all zone cluster nodes, install the HA for Samba data service IPS package from the Oracle Solaris Cluster 4.2 repository.**

```
zc-node-N# pkg install --accept ha-cluster/data-service/samba
```

▼ Example: Configure the Logical Host and Storage Cluster Resources for Samba

1. **On one zone-cluster node, register the `SUNW.gds` and `SUNW.HAStoragePlus` resource types within the zone cluster.**

```
zc-node-1# clresourcetype register SUNW.gds SUNW.HAStoragePlus
```

2. **On one zone-cluster node, create the Samba resource group `samba-rg`.**

```
zc-node-1# clresourcegroup create samba-rg
```

3. **On one zone-cluster node, create the logical host resource `samba-lh-rs` in the `samba-rg` resource group.**

```
zc-node-1# clreslogicalhostname create -g samba-rg -h samba-lh samba-lh-rs
```

4. **On both zone-cluster nodes, create the `/failover/samba` directory.**

```
zc-node-N# mkdir -p /failover/samba
```

5. **On one zone-cluster node, create the HAStoragePlus resource `samba-hasp-rs` in the `samba-rg` resource group.**

```
zc-node-1# clresource create -t SUNW.HAStoragePlus -g samba-rg \
-p zpools=samba_zp samba-hasp-rs
```

6. **On one zone-cluster node, enable the `samba-rg` resource group.**

```
zc-node-1# clresourcegroup online -eM samba-rg
```

▼ Example: Create the Samba `smb.conf` Configuration File

The Samba files are located where the Samba shares and `smb.conf` files are stored. The HA for Samba data service requires that these files be stored within a configuration directory that reflects the NetBIOS name, which in this example is `samba-lh` for the Samba or winbind instance.

1. **Create the Samba configuration directory.**

```
zc-node-1# mkdir /failover/samba/samba-lh
zc-node-1# cd /failover/samba/samba-lh
zc-node-1# mkdir -p lib logs private shares/data var/locks
zc-node-1# chmod 1777 /failover/samba/samba-lh/shares/data
```

2. **Create the Samba `smb.conf` configuration file.**

For more information about possible configuration options, see the example configuration file `/etc/samba/smb.conf-example`, the `smb.conf(5)` man page, and the online documentation available at <https://www.samba.org>.

```
zc-node-1# vi /failover/samba/samba-lh/lib/smb.conf
```

```
[global]
workgroup = OSC
server string = Samba (%v) Server (%h)
security = ADS
log file = /failover/samba/samba-lh/logs/log.%m
max log size = 50
realm = OSC.EXAMPLE.COM
passdb backend = smbpasswd
interfaces = 192.168.0.10/24
wins support = yes
dns proxy = no
netbios name = samba-lh
bind interfaces only = yes
pid directory = /failover/samba/samba-lh/var/locks
private dir = /failover/samba/samba-lh/private
smb passwd file = /failover/samba/samba-lh/private/smbpasswd
lock dir = /failover/samba/samba-lh/var/locks
winbind cache time = 30
allow trusted domains = no
idmap config * : backend = tdb
idmap config * : range = 10001-20000
idmap config DOMAIN : backend = rid
idmap config DOMAIN : range = 10000-20000
idmap config DOMAIN : base_riid = 0
winbind enum groups = yes
winbind enum users = yes
winbind use default domain = yes

[homes]
comment = Home Directories
browseable = no
writable = yes

[scmondird]
comment = Monitor directory for Oracle Solaris Cluster
path = /tmp
browsable = no

[data]
path = /failover/samba/samba-lh/shares/data
comment = Data Share
read only = no
writable = yes
```

3. Verify the smb.conf configuration file for correctness.

```
zc-node-1# testparm -s /failover/samba/samba-lh/lib/smb.conf
```


▼ Example: Create the Kerberos, PAM, and Name Service Switch Configuration for winbind

1. On all zone-cluster nodes, create the Kerberos `krb5.conf` file.

```
zc-node-N# vi /etc/krb5/krb5.conf

[libdefaults]
    default_realm = OSC.EXAMPLE.COM

[realms]
    OSC.EXAMPLE.COM = {
        kdc = 192.168.0.20
        admin_server = 192.168.0.20
    }

[domain_realm]
    .your.domain.name = OSC.EXAMPLE.COM
    your.domain.name = OSC.EXAMPLE.COM

[logging]
    default = FILE:/var/krb5/kdc.log
    kdc = FILE:/var/krb5/kdc.log
    kdc_rotate = {
        period = 1d
        versions = 10
    }

[appdefaults]
    kinit = {
        renewable = true
        forwardable= true
    }
```

2. Ensure that the global cluster nodes have network time protocol (NTP) configured to be in sync with the time used by the Windows Active Directory server.

Compare output of the `ntpq -p` command on all physical cluster nodes and the equivalent on the Windows Active Directory server. Kerberos relies on synchronized time between the systems.

3. On all zone-cluster nodes, activate the PAM configuration file for winbind.

```
zc-node-N# cp -p /etc/pam.conf /etc/pam.conf.orig
zc-node-N# cp /etc/pam.conf-winbind /etc/pam.conf
```

4. **On all zone-cluster nodes, configure the name service switch to resolve user and group from winbind.**

```
zc-node-N# svccfg -s name-service/switch setprop config/password = \"files winbind\"
zc-node-N# svccfg -s name-service/switch setprop config/group = \"files winbind\"
zc-node-N# svcadm refresh name-service/switch
```

5. **On all zone-cluster nodes, disable the name service cache daemon.**

```
zc-node-N# svcadm disable name-service/cache
```

6. **From one zone-cluster node, join the Active Directory domain.**

```
zc-node-1# net -s /failover/samba/samba-lh/lib/smb.conf ADS JOIN -U Administrator
```

▼ Example: Verify the Samba Configuration

1. **Log in to zc-node-1.**

2. **Start and test the winbindd program.**

```
zc-node-1# winbindd -s /failover/samba/samba-lh/lib/smb.conf
zc-node-1# getent passwd
    You should also see the users that are configured within the Active Directory domain
zc-node-1# getent group
    You should also see the groups that are configured within the Active Directory domain
zc-node-1# getent passwd hasmb
    You should see the passwd entry for user "hasmb"
```

3. **Start and test the smbd and nmbd programs.**

Confirm that the command output contains no errors.

```
zc-node-1# smbd -s /failover/samba/samba-lh/lib/smb.conf -D
zc-node-1# nmbd -s /failover/samba/samba-lh/lib/smb.conf -D
zc-node-1# nmblookup samba-lh
zc-node-1# smbclient -s /failover/samba/samba-lh/lib/smb.conf '\\samba-lh\scmndir' \
-U hasmb -c 'pwd;exit'
zc-node-1# smbclient -s /failover/samba/samba-lh/lib/smb.conf -N -L samba-lh
```

4. **On a Windows system that is part of the same Active Directory domain, configure usage of the SMB share.**

- a. **Go to My Computer -> Tools -> Map Network drive.**

- b. Select a drive letter and set the `\\samba-lh\data` directory as the share.
 - c. Create and access files in the share.
5. Verify that you can confirm access from the Windows system that you configured in [Step 4](#).

```
zc-node-1# smbstatus -s /failover/samba/samba-lh/lib/smb.conf
```

6. Stop the manually started Samba components.

```
zc-node-1# pkill 'smbd|nmbd|winbindd'
```

7. Switch the `samba-rg` resource group to `zc-node-2`.

```
zc-node-1# clresourcegroup switch -n zc-node-2 samba-rg
```

8. Repeat this procedure on `zc-node-2`.

▼ Example: Enable Samba Software to Run in the Cluster

1. On all zone-cluster nodes, set up the encrypted Samba fault monitor user password.

```
zc-node-N# touch /var/cluster/scsmb_key /opt/SUNWscsmb/.samba-lh_passwd
zc-node-N# chmod 600 /var/cluster/scsmb_key /opt/SUNWscsmb/.samba-lh_passwd
```

```
zc-node-N# dd if=/dev/urandom of=/var/cluster/scsmb_key bs=16 count=1
zc-node-N# chmod 400 /var/cluster/scsmb_key
```

```
zc-node-N# echo 'password' | /usr/sfw/bin/openssl enc -aes128 -e -pass \
file:/var/cluster/scsmb_key -out /opt/SUNWscsmb/.samba-lh_passwd
zc-node-N# chmod 400 /opt/SUNWscsmb/.samba-lh_passwd
```

2. On all zone-cluster nodes, confirm that the decrypted password matches the password configured for user `hasmb` within the Windows Active Directory domain.

```
zc-node-N# /usr/sfw/bin/openssl enc -aes128 -d -pass file:/var/cluster/scsmb_key \
-in /opt/SUNWscsmb/.samba-lh_passwd
```

3. From one node of the zone cluster, create the HA for Samba configuration file for the `winbind` resource.

```
zc-node-1# vi /var/tmp/samba_config-winbind-rs
```

```
RS=winbind-rs
RG=samba-rg
RS_LH=samba-lh-rs
RS_HAS=samba-hasp-rs
SERVICES="winbindd"
```

```
BINDIR=/usr/bin
SBINDIR=/usr/sbin
CFGDIR=/failover/samba/samba-lh
FMUSER=hasmb
```

```
SAMBA_LOGDIR=/failover/samba/samba-lh/logs
SAMBA_FMPASS=encrypted
SAMBA_FMDOMAIN=OSC
```

```
WINBIND_DISCACHE=FALSE
WINBIND_SINGLEMODE=FALSE
```

```
RS_ZONE=
LHOST=
PROJECT=default
TIMEOUT=30
```

4. From one node of the zone cluster, create the HA for Samba configuration file for the Samba resource.

```
zc-node-1# vi /var/tmp/samba_config-samba-rs
```

```
RS=samba-rs
RG=samba-rg
RS_LH=samba-lh-rs
RS_HAS=samba-hasp-rs
SERVICES="smbd,nmbd"
```

```
BINDIR=/usr/bin
SBINDIR=/usr/sbin
CFGDIR=/failover/samba/samba-lh
FMUSER=hasmb
```

```
SAMBA_LOGDIR=/failover/samba/samba-lh/logs
SAMBA_FMPASS=encrypted
SAMBA_FMDOMAIN=OSC
```

```
WINBIND_DISCACHE=FALSE
WINBIND_SINGLEMODE=FALSE
```

```
RS_ZONE=
LHOST=
PROJECT=default
TIMEOUT=30
```

5. From one node of the zone cluster, register the HA for Samba data service for winbind and Samba.

```
zc-node-1# /opt/SUNWscsmb/util/samba_register -f /var/tmp/samba_config-winbind-rs
zc-node-1# /opt/SUNWscsmb/util/samba_register -f /var/tmp/samba_config-samba-rs
```

6. From one node of the zone cluster, configure the resource dependency between the samba-rs and winbind-rs resources.

```
zc-node-1# clresource set -p Resource_dependencies+=winbind-rs samba-rs
```

7. From one node of the zone cluster, enable the resources.

```
zc-node-1# clresource enable winbind-rs
zc-node-1# clresource enable samba-rs
```

▼ Example: Verify the HA for Samba Resource Group

Before You Begin Ensure that the samba-rg resource group is online on zone-cluster node zc-node-1.

1. Verify the status of the cluster resources and resource group.

```
zc-node-1# clresourcegroup status samba-rg
zc-node-1# clresource status -g samba-rg
```

2. Verify the Samba functionality *without* manually starting the Samba daemons.

Confirm that the command output contains no errors.

```
zc-node-1# getent passwd
    You should also see the users that are configured within the Active Directory domain
zc-node-1# getent group
    You should also see the groups that are configured within the Active Directory domain
zc-node-1# getent passwd hasmb
    You should see the passwd entry for user hasmb
zc-node-1# nmblookup samba-lh zc-node-1# smbclient \
-s /failover/samba/samba-lh/lib/smb.conf \
'\\samba-lh\scmondir' -U hasmb -c 'pwd;exit'
```

```
zc-node-1# smbclient -s /failover/samba/samba-lh/lib/smb.conf -N -L samba-lh
```

3. **On a Windows system that is part of the same Active Directory domain, configure usage of the SMB share.**
 - a. **Go to My Computer -> Tools -> Map Network drive.**
 - b. **Select a drive letter and set the \\samba-lh\data directory as the share.**
 - c. **Create and access files in the share.**
4. **Verify that you can confirm access from the Windows system that you configured in [Step 3](#).**

```
zc-node-1# smbstatus -s /failover/samba/samba-lh/lib/smb.conf
```

5. **Switch the samba-rg resource group to the other zone cluster node, zc-node-2, and perform [Step 2](#) and [Step 3](#).**

```
zc-node-1# clresourcegroup switch -n zc-node-2-samba-rg
```

Index

C

command

- clresourcetype command, 22
- cluster command, 30

configuration

- smb.conf configuration file, 18
- verifying HA for Samba, 30
- verifying Samba, 25

configuring

- Samba, 19

D

dependencies, 18

deployment examples

- Samba packaged with Oracle Solaris 11, 35

E

examples

- Samba packaged with Oracle Solaris 11, 35

F

fault monitor

- overview, 31
- probing algorithm, 31
- resource properties, 31

I

installing

HA for Samba, 27

O

Oracle Solaris Cluster software

publisher, 27, 27

overview

- installation, 13
- product, 13

P

package

HA for Samba, 27

publisher

- Oracle Solaris Cluster software, 27, 27
- Samba software, 21

R

requirements, 14

resource types, 22

restrictions, 14

S

Samba software

publisher, 21

software package

HA for Samba, 27

