



Sun Cluster Data Service for Kerberos Guide for Solaris OS



Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 819-5415-10
December 2006 Revision A

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

U.S. Government Rights – Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2006 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux États-Unis et dans d'autres pays.

Cette distribution peut comprendre des composants développés par des tierces personnes.

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des États-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.

Contents

Preface	5
1 Installing and Configuring Sun Cluster for Kerberos	9
Sun Cluster HA for Kerberos	9
Installing and Configuring Sun Cluster HA for Kerberos	11
Configuring Sun Cluster HA for Kerberos in Non-Global Zones	12
▼ How to Configure Sun Cluster HA for Kerberos in Non-Global Zones	12
Installing Kerberos	14
▼ How to Install Kerberos	14
Installing the Sun Cluster HA for Kerberos Packages	19
▼ How to Install the Sun Cluster HA for Kerberos Packages	19
Registering and Configuring Sun Cluster HA for Kerberos	21
▼ How to Register and Configure Sun Cluster HA for Kerberos	21
▼ How to Configure the HASToragePlus Resource Type	24
Tuning the Sun Cluster HA for Kerberos Fault Monitor	25
Operations by the Fault Monitor During a Probe	26
Verifying Sun Cluster HA for Kerberos Installation and Configuration	27
▼ How to Verify Sun Cluster HA for Kerberos Installation and Configuration	27
A Sun Cluster HA for Kerberos Extension	29
Index	31

Preface

Sun Cluster Data Service for Kerberos Guide for Solaris OS explains how to install and configure Sun™ Cluster HA for Kerberos on both SPARC® based systems and x86 based systems.

Note – This Sun Cluster release supports systems that use the SPARC and x86 families of processor architectures: UltraSPARC, SPARC64, and AMD64. In this document, the label x86 refers to systems that use the AMD64 family of processor architectures.

This document is intended for system administrators with extensive knowledge of Sun software and hardware. Do not use this document as a planning or presales guide. Before reading this document, you should have already determined your system requirements and purchased the appropriate equipment and software.

The instructions in this book assume knowledge of the Solaris™ Operating System (Solaris OS) and expertise with the volume-manager software that is used with Sun Cluster software.

Note – Sun Cluster software runs on two platforms, SPARC and x86. The information in this document pertains to both platforms unless otherwise specified in a special chapter, section, note, bulleted item, figure, table, or example.

Using UNIX Commands

This document contains information about commands that are specific to installing and configuring Sun Cluster data services. The document does *not* contain comprehensive information about basic UNIX® commands and procedures, such as shutting down the system, booting the system, and configuring devices. Information about basic UNIX commands and procedures is available from the following sources:

- Online documentation for the Solaris Operating System
- Solaris Operating System man pages
- Other software documentation that you received with your system

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<code>machine_name%</code>
C shell for superuser	<code>machine_name#</code>
Bourne shell and Korn shell	<code>\$</code>
Bourne shell and Korn shell for superuser	<code>#</code>

Related Documentation

Information about related Sun Cluster topics is available in the documentation that is listed in the following table. All Sun Cluster documentation is available at <http://docs.sun.com>.

Topic	Documentation
Data service administration	<i>Sun Cluster Data Services Planning and Administration Guide for Solaris OS</i> Individual data service guides
Concepts	<i>Sun Cluster Concepts Guide for Solaris OS</i>
Overview	<i>Sun Cluster Overview for Solaris OS</i>
Software installation	<i>Sun Cluster Software Installation Guide for Solaris OS</i>
System administration	<i>Sun Cluster System Administration Guide for Solaris OS</i>
Hardware administration	<i>Sun Cluster 3.1 - 3.2 Hardware Administration Manual for Solaris OS</i> Individual hardware administration guides
Data service development	<i>Sun Cluster Data Services Developer's Guide for Solaris OS</i>
Error messages	<i>Sun Cluster Error Messages Guide for Solaris OS</i>
Command and function reference	<i>Sun Cluster Reference Manual for Solaris OS</i>

For a complete list of Sun Cluster documentation, see the release notes for your release of Sun Cluster at <http://docs.sun.com>.

Related Third-Party Web Site References

Third-party URLs that are referenced in this document provide additional related information.

Note – Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- [Documentation](http://www.sun.com/documentation/) (<http://www.sun.com/documentation/>)
- [Support](http://www.sun.com/support/) (<http://www.sun.com/support/>)
- [Training](http://www.sun.com/training/) (<http://www.sun.com/training/>)

Getting Help

If you have problems installing or using Sun Cluster, contact your service provider and provide the following information:

- Your name and email address (if available)
- Your company name, address, and phone number
- The model number and serial number of your systems
- The release number of the Solaris Operating System (for example, Solaris 10)
- The release number of Sun Cluster (for example, Sun Cluster 3.2)

Use the following commands to gather information about each node on your system for your service provider.

Command	Function
<code>prtconf -v</code>	Displays the size of the system memory and reports information about peripheral devices
<code>psrinfo -v</code>	Displays information about processors
<code>showrev -p</code>	Reports which patches are installed
<code>SPARCprtdiag -v</code>	Displays system diagnostic information
<code>/usr/cluster/bin/clnode show-rev</code>	Displays Sun Cluster release and package version information

Also have available the contents of the `/var/adm/messages` file.

Installing and Configuring Sun Cluster for Kerberos

This chapter describes the steps to install and configure the Sun Cluster HA for Kerberos data service on Sun Cluster servers.

This chapter contains the following sections:

- “Sun Cluster HA for Kerberos” on page 9
- “Installing and Configuring Sun Cluster HA for Kerberos” on page 11
- “Configuring Sun Cluster HA for Kerberos in Non-Global Zones” on page 12
- “Installing Kerberos” on page 14
- “Installing the Sun Cluster HA for Kerberos Packages” on page 19
- “Registering and Configuring Sun Cluster HA for Kerberos” on page 21
- “Tuning the Sun Cluster HA for Kerberos Fault Monitor” on page 25
- “Verifying Sun Cluster HA for Kerberos Installation and Configuration” on page 27

Sun Cluster HA for Kerberos

You must configure Sun Cluster HA for Kerberos as a failover data service. For conceptual information about failover data services, see Chapter 1, “Planning for Sun Cluster Data Services,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* and the *Sun Cluster Concepts Guide for Solaris OS*.

Kerberos servers have two daemons:

<code>krb5kdc(1M)</code>	Authentication service
<code>kadmind(1M)</code>	Principal or policy administration service

The `krb5kdc` daemon runs on both master and slave Key Distribution Center (KDC) servers. This service provides redundancy because an environment can have a master and one or more slaves that are running this process.

The `kadmind` daemon runs only on the master server and can handle requests that make updates to the principal/policy database. This single point of failure makes update requests more fragile than `krb5kdc`. By clustering the master KDC in the Kerberos environment you can provide update requests with greater availability.

For an introduction to Kerberos concepts, refer to Part VI, “Kerberos Service,” in *System Administration Guide: Security Services*.

Figure 1-1 lists the Kerberos components of a Sun Cluster environment.

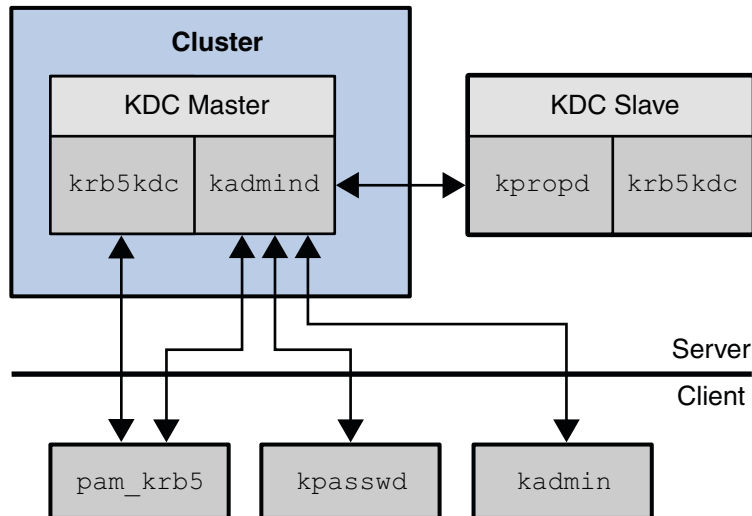


FIGURE 1-1 Kerberos Components in the Sun Cluster Environment

In Figure 1-1, `pam_krb5(5)`, `kpasswd(1)`, `kpropd(1M)`, and `kadmin(1M)` all send requests to `kadmind` directly. `pam_krb5` and `kpasswd` make update requests when changing a users password. `kadmin` is used for general administration of the principal and policy database.

Figure 1-2 shows how databases and configuration information are shared between the cluster nodes and zones through a global or failover file system.

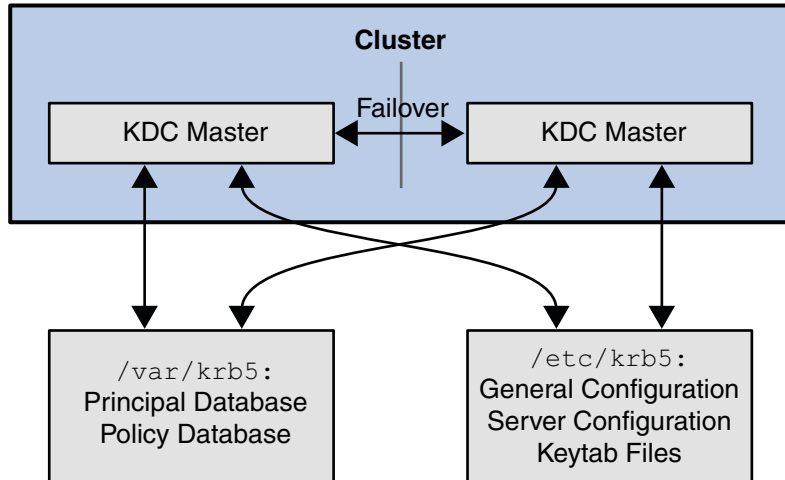


FIGURE 1-2 Database and Configuration Sharing

The configuration and keytab files are placed in `/etc/krb5`. The databases and logging files are kept under `/var/krb5`. By having these directories on a shared file system, you ensure that the database and configuration are identical. During failover, there should be little impact on client ticket requests, especially if there are slaves in the environment because slaves could be used to service client tickets during the failover period.

Clients that have already established sessions with `kadmin` by using the `kadmin` command are dropped after a failover on the cluster. Given the amount of privileges usually given for administrative principals, active `kadmin` sessions should not be left unattended. They should not run for an extended period of time. This means that `kadmin` session drops should not occur frequently because they are short lived processes.

Installing and Configuring Sun Cluster HA for Kerberos

Table 1-1 lists the tasks for installing and configuring Sun Cluster HA for Kerberos. Perform these tasks in the order in which they are listed unless otherwise indicated.

TABLE 1-1 Task Map: Installing and Configuring Sun Cluster HA for Kerberos

Task	Instructions
(Optional) Configure Sun Cluster HA for Kerberos in Non-Global Zones	“Configuring Sun Cluster HA for Kerberos in Non-Global Zones” on page 12
Install Kerberos	“Installing Kerberos” on page 14

TABLE 1-1 Task Map: Installing and Configuring Sun Cluster HA for Kerberos (Continued)

Task	Instructions
Install Sun Cluster HA for Kerberos packages	“Installing the Sun Cluster HA for Kerberos Packages” on page 19
Register and Configure Sun Cluster HA for Kerberos	“Registering and Configuring Sun Cluster HA for Kerberos” on page 21
(Optional) Tune the Sun Cluster HA for Kerberos fault monitor	“Tuning the Sun Cluster HA for Kerberos Fault Monitor” on page 25
Verify the Sun Cluster HA for Kerberos installation and configuration	“Verifying Sun Cluster HA for Kerberos Installation and Configuration” on page 27

Configuring Sun Cluster HA for Kerberos in Non-Global Zones

You can configure the Sun Cluster HA for Kerberos service within a non-global zone on Solaris 10 and later versions of the operating system. Given that all the realm's keys are stored in the KDC's principal database, it is helpful to compartmentalize access to system resources, such as file systems, into a non-global zone.

Note – Sun cluster software allows you to create different zones on the same node in which to deploy the Kerberos failover resources, but to provide high availability, create the zones deploying Kerberos failover resources on different nodes.

Note – Kerberos data service is supported on a sparse root non-global zone.

▼ How to Configure Sun Cluster HA for Kerberos in Non-Global Zones

Perform this procedure only if you want to configure the Sun Cluster HA for Kerberos service within a non-global zone.

Note – Configuring the Sun Cluster HA for Kerberos service in a global zone is similar to [“Installing Kerberos” on page 14](#) on a node.

If you do not want to configure the Sun Cluster HA for Kerberos service within a non-global zone, do not perform this procedure. Instead, go to [“Installing Kerberos” on page 14](#).

This procedure is written for use on a global file system. In this procedure, the following parameters are used:

- Global zone: global
- Non-global zone: sparse_zone
- Global file system: /global/fs

► **Create the non-global zone directory and mount it from the global zone. Perform this on each of the cluster nodes.**

```
sparse_zone# mkdir -p /global/fs
```

```
global# zonecfg -z sparse_zone
```

```
zonecfg:sparse_zone> add fs
zonecfg:sparse_zone:fs> set dir=/global/fs
zonecfg:sparse_zone:fs> set special=/global/fs
zonecfg:sparse_zone:fs> set type=lofs
zonecfg:sparse_zone:fs> end
zonecfg:sparse_zone> verify
zonecfg:sparse_zone> commit
zonecfg:sparse_zone> exit
```

```
global# zoneadm -z sparse_zone reboot
```

Where */global/fs* is a global file system that has already been configured in the global zone.

Note – The non-global zone's path must be identical to the path of the global zone.

Note – To simplify cluster administration, use the same non-global zone name on each node, where resource groups are to be brought online in the non-global zone.

Next Steps When you have configured the file system on all the non-global zones, go to [“How to Install Kerberos” on page 14](#). Perform the steps in that procedure in the non-global rather than the global zone.

Installing Kerberos

This section describes the steps to install Kerberos and to enable Kerberos to run as Sun Cluster HA for Kerberos.

Sun Cluster HA for Kerberos uses the Kerberos server and mechanism libraries co-packaged with the Solaris 10 operating system or later versions of the operating system. See the `krb5.conf(4)` and `kdc.conf(4)` man pages for information on how to configure the Kerberos environment. The Sun Cluster configuration for Kerberos differs from the Solaris configuration for Kerberos in the following ways:

- The Kerberos principal and policy databases are located on the cluster file system, not on a local file system. “[How to Install Kerberos](#)” on page 14 describes how to configure the server by using a global file system. However, the server can be configured with the `HAStoragePlus` file system if your environment is heavily loaded with write requests.
- A relocatable IP address, not the name of a physical host, identifies the name of a Kerberos server.

▼ How to Install Kerberos

In this procedure, the following parameters are used:

- Realm name = `EXAMPLE.COM`
- DNS domain name = `example.com`
- Cluster physical node names = `pkdc1.example.com` and `pkdc2.example.com`
- Cluster logical hostname = `kdc-1.example.com`

1 Become superuser on a cluster member.

2 Choose the logical hostname that will provide the Kerberos service.

Select the logical hostname so that it corresponds to an IP address set up when you installed the Sun Cluster software. See the *Sun Cluster Concepts Guide for Solaris OS* for details about logical hostnames.

3 Create the `krb5.conf`, `kdc.conf`, and the other configuration files required to run a Kerberos server, then run the command `kdb5_util(1M)` as described in the Chapter 23, “Configuring the Kerberos Service (Tasks),” in *System Administration Guide: Security Services*.

When populating the hostnames in these configuration files, ensure that they refer to the host's logical name, not the physical name.

Note – This detail ensures that applications running in the same zone as the logical hostname are configured to the corresponding IP addresses.

Here is an example of configuration files with the logical hostnames:

```
pkdc1# cat /etc/krb5/krb5.conf
```

```
[libdefaults]
    default_realm = EXAMPLE.COM

[realms]
    EXAMPLE.COM = {
        kdc = kdc-1.example.com
        admin_server = kdc-1.example.com
    }

[domain_realm]
    .example.com = 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
    }
```

```
pkdc1# cat /etc/krb5/kdc.conf
```

```
[kdcdefaults]
    kdc_ports = 88,750

[realms]
    ACME.COM = {
        profile = /etc/krb5/krb5.conf
        database_name = /var/krb5/principal
        admin_keytab = /etc/krb5/kadm5.keytab
        acl_file = /etc/krb5/kadm5.acl
        kadmind_port = 749
        max_life = 8h 0m 0s
        max_renewable_life = 7d 0h 0m 0s
        default_principal_flags = +preauth
    }
```

Make sure that you also have a valid `/etc/resolv.conf` file and `/etc/nsswitch.conf` file configured, for example:

```
pkdc1# cat /etc/resolv.conf
```

```
domain example.com
```

```
nameserver 1.2.3.4
```

```
nameserver 1.2.3.5
```

```
pkdc1# grep dns nsswitch.conf
```

```
hosts:          files nis dns
```

```
ipnodes:        files nis dns
```

4 Create the KDC database by running the `kdb5_util(1M)`

```
pkdc1# kdb5_util create
```

```
Initializing database '/var/krb5/principal' for realm 'EXAMPLE.COM',  
master key name 'K/M@EXAMPLE.COM'
```

```
You will be prompted for the database Master Password.
```

```
It is important that you NOT FORGET this password.
```

```
Enter KDC database master key:<Type the new master key password>
```

```
Re-enter KDC database master key:<Type the above new master key password>
```

5 Add the following line in the `/etc/krb5/kadm5.acl` file:

```
sckrb5-probe/admin@EXAMPLE.COM i
```

Where:

EXAMPLE.COM Realm name chosen in [Step 3](#)

i The privilege that enables queries to the database for the
sckrb5-probe/admin principal

6 Start the `kadmin.local` command.

```
pkdc1# kadmin.local
```

Authenticating as principal `host/admin@EXAMPLE.COM` with password

a. Use the `kadmin.local` command to add `kadmin` and `changepw` service principals for the fully qualified logical hostname for the cluster, `kdc-1.example.com`.

```
kadmin.local: ank -randkey -allow_tgs_req kadmin/kdc-1.example.com
```

NOTICE: no policy specified for `kadmin/kdc-1.example.com@EXAMPLE.COM`; assigning "default" Principal "`kadmin/kdc-1.example.com@EXAMPLE.COM`" created.

```
kadmin.local: ank -randkey -allow_tgs_req +password_changing_service \
changepw/kdc-1.example.com
```

NOTICE: no policy specified for `changepw/kdc-1.example.com@EXAMPLE.COM`; assigning "default" Principal "`changepw/kdc-1.example.com@EXAMPLE.COM`" created.

```
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab kadmin/kdc-1.example.com changepw/kdc-1.example.com
Entry for principal kadmin/kdc-1.example.com with kvno 3, encryption type AES-+ 128 CTS mode with \
96-bit SHA-1 HMAC added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal kadmin/kdc-1.example.com with kvno 3, encryption type Triple
DES cbc mode with HMAC/sha1 added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal kadmin/kdc-1.example.com with kvno 3, encryption type
ArcFour with HMAC/md5 added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal kadmin/kdc-1.example.com with kvno 3, encryption type
DES cbc mode with RSA-MD5 added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal changepw/kdc-1.example.com with kvno 3, encryption type
AES-128 CTS mode with 96-bit SHA-1 HMAC added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal changepw/kdc-1.example.com with kvno 3, encryption type
Triple DES cbc mode with HMAC/sha1 added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal changepw/kdc-1.example.com with kvno 3, encryption type
ArcFour with HMAC/md5 added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal changepw/kdc-1.example.com with kvno 3, encryption type
DES cbc mode with RSA-MD5 added to keytab WRFILE:/etc/krb5/kadm5.keytab.
```

b. Add the new service principals for the host services for the fully qualified logical hostname for the cluster, `kdc-1.example.com`:

```
kadmin.local: ank -randkey host/kdc-1.example.com
```

NOTICE: no policy specified for `host/kdc-1.example.com@EXAMPLE.COM`; assigning "default" Principal "`host/kdc-1.example.com@EXAMPLE.COM`" created.

```
kadmin.local: ktadd host/kdc-1.example.com
```

```
Entry for principal host/kdc-1.example.com with kvno 3, encryption type AES-128 CTS mode with 96-bit SHA-1 \
```

HMAC added to keytab WRFILE:/etc/krb5/krb5.keytab.

Entry for principal host/kdc-1.example.com with kvno 3, encryption type Triple DES cbc mode with HMAC/sha1 \ added to keytab WRFILE:/etc/krb5/krb5.keytab.

Entry for principal host/kdc-1.example.com with kvno 3, encryption type ArcFour with HMAC/md5 added to keytab \ WRFILE:/etc/krb5/krb5.keytab.

Entry for principal host/kdc-1.example.com with kvno 3, encryption type DES cbc mode with RSA-MD5 added to \ keytab WRFILE:/etc/krb5/krb5.keytab.

kdc-1.example.com Fully qualified logical hostname for the cluster

c. Add a new service principal for the kiprof service for the fully qualified logical hostname for the cluster, kdc-1.example.com.

kadmin.local: **ank -randkey kiprof/kdc-1.example.com**

NOTICE: no policy specified for kiprof/kdc-1.example.com@EXAMPLE.COM; assigning "default"

Principal "kiprof/kdc-1.example.com@EXAMPLE.COM" created.

kadmin.local: **ktadd -k /etc/krb5/kadm5.keytab kiprof/kdc-1.example.com**

Entry for principal kiprof/kdc-1.example.com with kvno 3, encryption type AES-128 CTS mode with 96-bit \ SHA-1 HMAC added to keytab WRFILE:/etc/krb5/kadm5.keytab.

Entry for principal kiprof/kdc-1.example.com with kvno 3, encryption type Triple DES cbc mode with HMAC/sha1 \ added to keytab WRFILE:/etc/krb5/kadm5.keytab.

Entry for principal kiprof/kdc-1.example.com with kvno 3, encryption type ArcFour with HMAC/md5 added to \ keytab WRFILE:/etc/krb5/kadm5.keytab.

Entry for principal kiprof/kdc-1.example.com with kvno 3, encryption type DES cbc mode with RSA-MD5 added \ to keytab WRFILE:/etc/krb5/kadm5.keytab.

7 Move the /etc/krb5 and /var/krb5 directories to either a global or a failover file system.

For example, move /etc/krb5 and /var/krb5 to a global file system, /global/fs/, as follows:

```
pkdc1# mv /etc/krb5 /global/fs/krb-conf
```

```
pkdc1# mv /var/krb5 /global/fs/krb-db
```

See the *Sun Cluster Software Installation Guide for Solaris OS* for information on setting up cluster file systems.

8 Create symbolic links back to the /etc/krb5 and /var/krb5 directories:

```
pkdc1# ln -s /global/fs/krb-conf /etc/krb5
```

```
pkdc1# ln -s /global/fs/krb-db /var/krb5
```

- 9 Repeat the symbolic link creation on all the other cluster nodes or zones.

```
pkdc2# mv /etc/krb5 /etc/krb5.old
```

```
pkdc2# mv /var/krb5 /var/krb5.old
```

```
pkdc2# ln -s /global/fs/krb-conf /etc/krb5
```

```
pkdc2# ln -s /global/fs/krb-db /var/krb5
```

Installing the Sun Cluster HA for Kerberos Packages

If you did not install the Sun Cluster HA for Kerberos packages during your initial Sun Cluster installation, perform this procedure to install the packages. To install the packages, use the Sun Java™ Enterprise System Installation Wizard.

▼ How to Install the Sun Cluster HA for Kerberos Packages

Perform this procedure on each cluster node where you are installing the Sun Cluster HA for Kerberos packages.

You can run the Sun Java Enterprise System Installation Wizard with a command-line interface (CLI) or with a graphical user interface (GUI). The content and sequence of instructions in the CLI and the GUI are similar.

Note – Even if you plan to configure this data service to run in non-global zones, install the packages for this data service in the global zone. The packages are propagated to any existing non-global zones and to any non-global zones that are created after you install the packages.

Before You Begin Ensure that you have the Sun Java Availability Suite DVD-ROM.

If you intend to run the Sun Java Enterprise System Installation Wizard with a GUI, ensure that your DISPLAY environment variable is set.

- 1 On the cluster node where you are installing the data service packages, become superuser.

2 Load the Sun Java Availability Suite DVD-ROM into the DVD-ROM drive.

If the Volume Management daemon `volm(1M)` is running and configured to manage DVD-ROM devices, the daemon automatically mounts the DVD-ROM on the `/cdrom` directory.

3 Change to the Sun Java Enterprise System Installation Wizard directory of the DVD-ROM.

- **If you are installing the data service packages on the SPARC® platform, type the following command:**

```
# cd /cdrom/dcdrom0/Solaris_sparc
```

- **If you are installing the data service packages on the x86 platform, type the following command:**

```
# cd /cdrom/dcdrom0/Solaris_x86
```

4 Start the Sun Java Enterprise System Installation Wizard.

```
# ./installer
```

5 When you are prompted, accept the license agreement.

If any Sun Java Enterprise System components are installed, you are prompted to select whether to upgrade the components or install new software.

6 From the list of Sun Cluster agents under Availability Services, select the data service for Kerberos.

7 If you require support for languages other than English, select the option to install multilingual packages.

English language support is always installed.

8 When prompted whether to configure the data service now or later, choose Configure Later.

Choose Configure Later to perform the configuration after the installation.

9 Follow the instructions on the screen to install the data service packages on the node.

The Sun Java Enterprise System Installation Wizard displays the status of the installation. When the installation is complete, the wizard displays an installation summary and the installation logs.

10 (GUI only) If you do not want to register the product and receive product updates, deselect the Product Registration option.

The Product Registration option is not available with the CLI. If you are running the Sun Java Enterprise System Installation Wizard with the CLI, omit this step

- 11 Exit the Sun Java Enterprise System Installation Wizard.
- 12 Unload the Sun Java Availability Suite DVD-ROM from the DVD-ROM drive.
 - a. To ensure that the DVD-ROM is not being used, change to a directory that does *not* reside on the DVD-ROM.
 - b. Eject the DVD-ROM.


```
# eject cdrom
```

Next Steps See “[Registering and Configuring Sun Cluster HA for Kerberos](#)” on page 21 to register Sun Cluster HA for Kerberos and to configure the cluster for the data service.

Registering and Configuring Sun Cluster HA for Kerberos

This section describes how to register and configure Sun Cluster HA for Kerberos.

▼ How to Register and Configure Sun Cluster HA for Kerberos

- Before You Begin** To perform this procedure, you need the following information about your configuration.
- The name of the resource type for Sun Cluster HA for Kerberos. This name is `SUNW.krb5`.
 - The names of the cluster nodes and the non-global zones on the nodes that master the data service.
 - The network resource that clients use to access the data service. You normally set up this IP address when you install the cluster. See the *Sun Cluster Concepts Guide for Solaris OS* document for details on network resources.

1 **Become superuser on a cluster node.**

2 **Register the resource type for the data service.**

```
# clresourcetype register SUNW.krb5
```

3 **Create a resource group for the network and Kerberos resources to use.**

```
# clresourcegroup create [-n node[,...]] resource-group
```

`-n node[,...]` Specifies an optional comma-separated list of zones that can master this resource group. Each entry in this list has the format `node`. Where `node` is the node name and address and `zone` specifies the name of a non-global Solaris

zone. To specify the global zone, or to specify a node without local zones, specify only *node*. These are the nodes or zones on which the data service can run. The order here determines the order in which the nodes or zones are considered as primary during failover. If all of the cluster nodes or zones are potential masters, you do not need to use the *-n* option.

This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

4 Verify that all of the network resources that are to be used have been added to your name service database.

You should have performed this verification during the Sun Cluster installation. See the Chapter 1, “Planning the Sun Cluster Configuration,” in *Sun Cluster Software Installation Guide for Solaris OS* for details.

Note – To avoid any failures because of name service lookup, verify that all of the network resources are present in the server's and client's `/etc/inet/hosts` file. Configure name service mapping in the `/etc/nsswitch.conf` file on the servers to first check the local files before trying to access NIS or NIS+.

5 Add a logical hostname to a resource group.

```
# clreslogicalhostname create -g resource-group \
-h logical-hostname,[logical-hostname] \
[-N netif@node[,...]] lhresource
```

- g *resource-group* Specifies the name of the resource group. This name can be your choice but must be unique for a resource group within the cluster.
- h *logical-hostname* Specifies a comma-separated list of network resources (logical hostname or shared address).
- N *netif@node[,...]* Specifies an optional, comma-separated list that identifies the IP Networking Multipathing groups that are on each node. *netif* can be given as an IP Networking Multipathing group name, such as `sc_ipmp0`. The node can be identified by the node name or node ID, such as `sc_ipmp0@1` or `sc_ipmp0@phys-schost-1`. If you do not specify *-N*, the `clreslogicalhostname` command attempts to set the `NetIFList` property for you based on available IPMP groups or public adapters and the subnet associated with the `HostnameList` property.
- lhresource* Specifies the logical hostname resource to be created in the associated resource group.

Note – If you require a fully qualified hostname, you must specify the fully qualified name with the `-h` option and you cannot use the fully qualified form in the resource name.

Note – Sun Cluster does not currently support the use of adapter names for *netif*.

6 Add a Kerberos application resource to the resource group.

```
# clresource create -g resource-group -t SUNW.krb5 \
[-p Network_resources_used=network-resource, ...] \
[-p Port_list=port-number/protocol] resource
```

`-p Network_resources_used=network-resource, ...`

Specifies a comma-separated list of network resources (logical hostnames or shared addresses) that Kerberos will use. If you do not specify this property, the value defaults to all of the network resources that are contained in the resource group.

`-p Port_list=port-number/protocol`

Specifies a port number and the protocol to be used. If you do not specify this property, the value defaults to `88/tcp, 749/tcp, 88/udp`.

`-t SUNW.krb5`

Specifies the name of the resource type to which this resource belongs. This entry is required.

`resource`

Specifies the name of the resource to be associated with the resource type `SUNW.krb5`.

The resource is created in the enabled state.

7 Bring the resource group online:

```
# clresourcegroup online -M resource-group
```

Example 1–1 Registering Failover Sun Cluster HA for Kerberos

The following example shows how to register Sun Cluster HA for Kerberos on a two-node cluster. At the end of this example, the `clresourcegroup` command starts Sun Cluster HA for Kerberos.

This example uses the following configuration parameters:

Cluster physical node names `pkdc1.example.com` and `pkdc2.example.com:sparse_zone`

Note – Kerberos is hosted in the global zone on `pkdc1.example.com` and in the non-global zone “`sparse_zone`” on `pkdc2.example.com`.

Cluster logical hostname	<code>kdc-1.example.com</code>
Resource group	<code>krb-rg</code> (for all of the resources)
Resources	<code>kdc-1</code> (logical hostname) and <code>krb-rs</code> (Kerberos application resource)

1. Register the Kerberos resource type.

```
# clresourcetype register SUNW.krb5
```

2. Create the resource group to contain all of the resources.

```
# clresourcegroup create -n pkdc1.example.com, pkdc2.example.com:sparse_zone krb-rg
```

3. Add the logical hostname resource to the resource group.

```
# clreslogicalhostname create -g krb-rg -h kdc-1 kdc-1
```

4. Add a Kerberos application resource to the resource group.

```
# clresource create -g krb-rg -t SUNW.krb5 krb-rs
```

5. Bring the failover resource group online.

```
# clresourcegroup online -M krb-rg
```

▼ How to Configure the HASStoragePlus Resource Type

This procedure describes how to configure the HASStoragePlus resource type. This resource type synchronizes actions between HASStorage and Sun Cluster HA for Kerberos and enables you to use a highly available local file system. It is, however, recommended that you use a global file system rather than using HASStoragePlus because Sun Cluster HA for Kerberos is not disk-intensive in most environments.

See “Relationship Between Resource Groups and Device Groups” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for background information.

This procedure uses the following configuration parameters:

- Cluster physical node names = `pkdc1.example.com` and `pkdc2.example.com:sparse_zone`
- Cluster logical hostname = `kdc-1.example.com`

- Resource group = krb-rg
- Kerberos application resource = krb-rs
- HAStoragePlus resource = krb-hasp-rs
- Logical hostname resource = kdc-1
- Device group associated with the file system:/global/dg1

Note – The /global/dg1 file system contains the krb-db and krb-conf directories which have symbolic links that point to /var/krb5 and /etc/krb5 respectively.

1 Register the Kerberos resource type.

```
# clresourcetype register SUNW.krb5
```

2 Create a resource group.

```
# clresourcegroup create -n pkdc1.example.com, pkdc2.example.com:sparse_zone krb-rg
```

3 Add the logical hostname resource to the resource group.

```
# clreslogicalhostname create -g krb-rg -h kdc-1
```

4 Add the Kerberos application resource to the resource group.

```
# clresource create -g krb-rg -t SUNW.krb5 krb-rs
```

5 Register the HAStoragePlus resource type

```
# clresourcetype register SUNW.HAStoragePlus
```

6 Add the HAStoragePlus resource to the resource group.

```
# clresource create -g krb-rg -t SUNW.HAStoragePlus \
-p FilesystemMounPoints=/global/dg1 \
-p AffinityOn=TRUE krb-hasp-rs
```

7 Bring the failover resource group online.

```
# clresourcegroup online -M krb-rg
```

Tuning the Sun Cluster HA for Kerberos Fault Monitor

The Sun Cluster HA for Kerberos fault monitor is contained in the resource that represents Kerberos. You create this resource when you register and configure Sun Cluster HA for Kerberos. For more information, see [“Registering and Configuring Sun Cluster HA for Kerberos” on page 21](#).

System properties and extension properties of this resource control the behavior of the fault monitor. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Sun Cluster installations. Therefore, you should tune the Sun Cluster HA for Kerberos fault monitor *only* if you need to modify this preset behavior.

Tuning the Sun Cluster HA for Kerberos fault monitor involves the following tasks:

- Setting the interval between fault monitor probes
- Setting the timeout for fault monitor probes
- Defining the criteria for persistent faults
- Specifying the failover behavior of a resource

Perform these tasks when you register and configure Sun Cluster HA for Kerberos. For more information, see the following sections:

- [“Registering and Configuring Sun Cluster HA for Kerberos” on page 21](#)
- [“Tuning Fault Monitors for Sun Cluster Data Services” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*](#)

The information that you need to tune the [“Registering and Configuring Sun Cluster HA for Kerberos” on page 21](#) fault monitor is provided in the follow subsection.

Operations by the Fault Monitor During a Probe

The probing consists of checking to see if `kadmind(1M)` and `krb5kdc(1M)` are listening to their respective ports. During more thorough probing a new principal is created, the principal is authenticated, and then this principal fetches itself from the database to test the administrative daemon, `kadmind`.

The probe executes the following steps.

1. Probe the ports for `kadmind(1M)` and `krb5kdc(1M)` to make sure that they are listening. Run the probe command by using the time-out value that the resource property `Probe_timeout` specifies. The probe is run every `Cheap_probe_interval`, which by default is every 30 seconds.
2. Every `Thorough_probe_interval` (by default 300 seconds) `kadmin.local(1M)` is used to add a principal. The probe then performs a `kinit(1)` with the newly created principal. The probe uses the newly created principal to run `kadmin(1M)` to retrieve its record from the principal database.
3. The result of these probe commands can be either fail or succeed. If Kerberos successfully responds, the probe returns to its infinite loop, waiting for the next probe time.

If the probe fails, the probe considers this scenario a failure of the Kerberos data service and records the failure in its history. The Kerberos probe considers every failure a complete failure.

4. Based on the success or failure history, a failure can cause a local restart or a data service failover. Refer to “Tuning Fault Monitors for Sun Cluster Data Services” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS* for further details.

Verifying Sun Cluster HA for Kerberos Installation and Configuration

Verify that you have correctly installed and configured Sun Cluster HA for Kerberos.

▼ How to Verify Sun Cluster HA for Kerberos Installation and Configuration

- 1 **Configure a Kerberos client to authenticate the newly created server as described in “Configuring Kerberos Clients” in *System Administration Guide: Security Services*.**

When referencing the server in the client’s configuration file, `/etc/krb5/krb5.conf`, specify the logical hostname of the server. An example of this name could be `kdc-1.example.com`.

- 2 **After the client is configured, test the authentication of a user principal by using `kinit(1)`.**

```
# kinit user_name
```

Password for `user_name@realm_name`:

`user_name` A user principal that you created previously

`realm_name` Indicates the realm name that was previously configured in the Kerberos environment.

You return to the shell prompt without any error message being output to the terminal.

- 3 **Test the Kerberos administration service.**

```
# kadmin -p administrative_principal_name
```

Authenticating as principal `administrative_principal_name@realm_name` with password.

Password for `administrative_principal_name@realm_name`:

```
kadmin: quit
```

`administrative_principal_name` An administrative principal that you had previously created.

realm_name Indicates the realm name that was previously configured in the Kerberos environment.

You return to the shell prompt without any error message being output to the terminal.

Sun Cluster HA for Kerberos Extension

You do not have to specify extension properties when creating a Kerberos resource. To configure extension properties when you create a resource, use the `-p` option of the `clresourcegroup(1CL)` command. To configure extension properties at a later stage, perform the procedures in Chapter 2, “Administering Data Service Resources,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*. For information about all of the Sun Cluster properties, see Appendix B, “Standard Properties,” in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

For details about system-defined properties, see the `r_properties(5)` man page and the `rg_properties(5)` man page.

The `SUNW.krb5` resource type represents the Kerberos application in a Sun Cluster configuration. The extension properties of this resource type are as follows:

Monitor_retry_count

Controls fault-monitor restarts. The property indicates the number of times that the process monitor facility restarts the fault monitor. The property corresponds to the `-n` option passed to the `pmfadm(1M)` command. The Resource Group Manager (RGM) counts the number of restarts in a specified time window. See the `Monitor_retry_interval` property for more information. Note that `Monitor_retry_count` refers to the restarts of the fault monitor itself, not to the Kerberos daemon.

Data type	Integer
Default	4
Range	Not applicable
Tunable	Anytime

Monitor_retry_interval

Indicates the time window in minutes during which the RGM counts fault-monitor failures. The property corresponds to the `-t` option passed to the `pmfadm(1M)` command. If the

number of times the fault monitor fails exceeds the value of the `Monitor_retry_count` property, the process monitor facility does not restart the fault monitor.

Data type Integer
Default 2 minutes
Range Not applicable
Tunable Anytime

Probe_timeout
Probe_timeout

Data type Integer
Default 90 seconds
Range Not applicable
Tunable Anytime

Index

C

`clnode` command, 8
commands, node information, 8

E

error messages, cluster, 8
extension properties, Sun Cluster HA for Kerberos, 29

F

failover data service, Sun Cluster HA for
Kerberos, 9-11
fault monitors
Sun Cluster HA for Kerberos
resource type for, 25
tuning, 25

G

global zone, 19

H

help, 8

I

installing
Kerberos, 14-19
Sun Cluster HA for Kerberos, 19-21
introduction, Sun Cluster HA for Kerberos, 9-11

K

`kadmind`, administration service, 9
Kerberos, installing, 14-19
Kerberos application, fault monitor, 25
`krb5kdc`, authentication service, 9

L

local zones, *See* non-global zones

M

messages file, 8
`Monitor_retry_count` extension property, `SUNW.krb5`
resource type, 29
`Monitor_retry_interval` extension property,
`SUNW.krb5` resource type, 29

N

non-global zones, 19

P

- packages, 19-21
- Probe_timeout extension property
 - SUNW.krb5 resource type, 30
- prtconf -v command, 8
- prtdiag -v command, 8
- psrinfo -v command, 8

R

- resource types, fault monitors, 25
- restrictions, zones, 19

S

- show-rev subcommand, 8
- showrev -p command, 8
- software packages, 19-21
- Sun Cluster HA for Kerberos
 - extension properties, 29
 - failover data service, 9-11
 - fault monitor, 25
 - installing, 19-21
 - introduction, 9-11
 - software packages
 - installing, 19-21
 - task map, 11-12
- system properties, effect on fault monitors, 26

T

- task map, Sun Cluster HA for Kerberos, 11-12
- technical support, 8
- tuning, fault monitors, 25

V

- /var/adm/messages file, 8

Z

- zones, 19