



Sun Cluster Geographic Edition Overview



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

Part No: 819-4243-10
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 Etats-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 Etats-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 Etats-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 Etats-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 Etats-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 CONTREFACON.

Contents

Preface	5
1 Introduction to Sun Cluster Geographic Edition Software	9
Business Continuity	9
Making Applications Highly Available With Sun Cluster Geographic Edition Software	10
Recovering From a Disaster	10
Key Features of Sun Cluster Geographic Edition Software	11
Administration and Configuration Tools	12
2 Key Concepts for Sun Cluster Geographic Edition	13
Data Replication	13
Replication Resource Groups	14
Device Groups	15
Cluster Partnerships	15
Protection Groups	16
Relationship Between Partnerships and Protection Groups	16
Protection Group Status	17
Application Resource Groups	18
Heartbeat Monitoring	18
Heartbeat Plug-ins	18
3 Sun Cluster Geographic Edition Architecture	19
Sun Cluster Geographic Edition Software Environment	19
Sun Cluster Geographic Edition Hardware Environment	21
Data Replication Configuration	21
Geographically Distributed Cluster Topology	23

Glossary27

Index31

Preface

Sun Cluster Geographic Edition Overview introduces the Sun™ Cluster Geographic Edition software by explaining the purpose of the product and the means by which Sun Cluster Geographic Edition achieves this purpose. This book also explains key concepts for Sun Cluster Geographic Edition. This document contains information about Sun Cluster Geographic Edition features and functionality.

Using UNIX Commands

This document contains information about commands that are used to install, configure, or administer a Sun Cluster Geographic Edition configuration. This document might not contain complete information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information:

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

Related Documentation

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

Topic	Documentation
Overview	<i>Sun Cluster Geographic Edition Overview</i>
Glossary	<i>Sun Java Enterprise System Glossary</i>
Hardware administration	Individual hardware administration guides
Software installation	<i>Sun Cluster Geographic Edition Installation Guide</i>

Topic	Documentation
System administration	<i>Sun Cluster Geographic Edition System Administration Guide</i> <i>Sun Cluster Geographic Edition Data Replication Guide for Sun StorEdge Availability Suite</i> <i>Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy</i> <i>Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility</i>
Command and function references	<i>Sun Cluster Geographic Edition Reference Manual</i>

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

Documentation, Support, and Training

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

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

Getting Help

If you have problems installing or using the Sun Cluster Geographic Edition system, 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 and serial numbers of your systems
- The release number of the operating system (for example, Solaris 9)
- The release number of the Sun Cluster Geographic Edition software (for example, 3.1 2006 Q4)

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>prtdiag -v</code>	Displays system diagnostic information
<code>scinstall -pv</code>	Displays Sun Cluster Geographic Edition software release and package version information
<code>scstat</code>	Provides a snapshot of the cluster status
<code>scconf -p</code>	Lists cluster configuration information
<code>scrgadm -p</code>	Displays information about installed resources, resource groups, and resource types
<code>geoadm status</code>	Displays runtime status of the local cluster

Also have available the contents of the `/var/adm/messages` file and the log files in `/var/opt/SUNWcacao/logs`.

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
<code>AaBbCc123</code>	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 <code>rm filename</code> .

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . A <i>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	machine_name%
C shell for superuser	machine_name#
Bourne shell and Korn shell	\$
Bourne shell and Korn shell for superuser	#

Introduction to Sun Cluster Geographic Edition Software

Sun Cluster Geographic Edition software is a layered extension of the Sun Cluster software. The Sun Cluster Geographic Edition software protects applications from unexpected disruptions by using multiple clusters that are separated by long distances and by using a redundant infrastructure that replicates data between these clusters. Data replication software enables applications that are running on a Sun Cluster Geographic Edition cluster to tolerate disasters by migrating services to a geographically separated secondary cluster. A disaster such as an earthquake, a fire, or a storm might disable the cluster at the primary site. If a disaster occurs, the Sun Cluster Geographic Edition cluster can continue to provide services by using the following levels of redundancy:

- A secondary cluster
- Duplicated application configuration on the secondary cluster
- Replicated data on the secondary cluster

This chapter gives a high-level overview of the Sun Cluster Geographic Edition product.

This chapter contains the following sections:

- “Business Continuity” on page 9
- “Making Applications Highly Available With Sun Cluster Geographic Edition Software” on page 10
- “Recovering From a Disaster” on page 10
- “Key Features of Sun Cluster Geographic Edition Software” on page 11
- “Administration and Configuration Tools” on page 12

Business Continuity

Business continuity is a wide-ranging subject that analyzes all aspects of how a business maintains service to its customers when faced with an unexpected disaster. When creating a business continuity plan, companies must trade off the cost of the additional required infrastructure such as hardware, software, telecommunications, and buildings, against the risks, such as the costs of a prolonged outage. As a result, systems that are critical to the business and those for which there is a legal requirement are the top priority.

For a service to be available, all the constituent pieces must also be available. The key issue is rapid recovery of individual service elements. Outages can occur from hardware failure such as component or power failures, or from software failures such as operating system panics and application crashes. Network connectivity failures can also affect service availability. Most of these failures can be masked through components redundancy or by having a standby server ready to take over the workload. The Sun Cluster Geographic Edition software is a building block for disaster tolerance which provides a framework that enables data services to be moved between a primary cluster and a geographically separated secondary cluster in a controlled fashion.

Making Applications Highly Available With Sun Cluster Geographic Edition Software

The Sun Cluster Geographic Edition software provides a suite of tools to manage and configure geographically separated clusters with a migration of services between sites. The Sun Cluster Geographic Edition software can manage availability across multiple physical locations through robust security, application service migration, and data replication to tolerate disaster across an enterprise system.

The Sun Cluster Geographic Edition product enables an improved combination of performance, cost, and separation of data recovery points. This combination contrasts with campus or metro clustering, which consists of a single cluster with widely separated nodes. The Sun Cluster Geographic Edition product provides the management and configuration tools for geographically separated clusters.

A configuration that is running the Sun Cluster Geographic Edition software consists of a set of clusters that are geographically distributed. The primary cluster provide application services, and the secondary cluster in the set is an alternative site that can take over the primary cluster services if a disaster occurs. The Sun Cluster Geographic Edition software manages configuration, data replication, and heartbeat monitoring between the two clusters and enables data to be decentralized across multiple disaster recovery sites.

Recovering From a Disaster

Disaster tolerance is the ability of a system to restore an application on a secondary cluster when the primary cluster fails. Disaster tolerance is based on data replication and failover. The Sun Cluster Geographic Edition software enables disaster tolerance by redundantly deploying the following:

- Highly available clusters that are geographically separated
- Data replication at either the host or the storage level
- Backups and restoration and data vaulting

Data replication is the process of continuously copying data from the primary cluster to the secondary cluster. Through data replication, the secondary cluster has a recent copy of the data on the primary cluster. The secondary cluster can be geographically separated from the primary cluster.

Failover is the automatic relocation of a resource group or device group from a primary cluster to a secondary cluster. If the primary cluster fails, the application and the data are immediately available on the secondary cluster.

The Sun Cluster Geographic Edition software supports two types of migration of services: a switchover and a takeover. A switchover is a planned migration of services from the primary cluster to the secondary cluster. During a switchover, the primary cluster is connected to the secondary cluster and coordinates the migration of services with the secondary cluster. This coordination enables the data replication to complete and ensures that services can be transferred from the primary cluster to the secondary cluster with minimal loss or corruption of data.

A takeover is an emergency migration of services from the primary cluster to the secondary cluster. A system administrator can initiate a takeover to recover from a disaster. Unlike a switchover, the primary cluster is not connected to the secondary cluster during a takeover. Therefore, the primary cluster cannot coordinate with the secondary cluster to migrate the services. Because of this lack of coordination, the risk of data loss and data corruption in a takeover is higher than it is with a switchover. The Sun Cluster Geographic Edition software uses dedicated recovery procedures during a takeover to minimize data loss and data corruption.

Key Features of Sun Cluster Geographic Edition Software

The Sun Cluster Geographic Edition product provides the following features:

- Failure detection of multiple clusters that are geographically separated
- Configurable heartbeat monitoring between clusters
- Application resource switchover from one cluster to another cluster
- Remote management of partner clusters through a graphical user interface (GUI) and a command-line interface (CLI)
- Data replication between geographically separated clusters
- Secure administration interfaces through role-based access control (RBAC)
- Secure Sockets Layer (SSL) authentication and encryption for communication between nodes or clusters
- Configurable IPsec security for data replication between clusters and for heartbeat communication between clusters
- Ability to automatically run a script when a heartbeat-loss notification is issued

The Sun Cluster Geographic Edition software provides tools for managing data replication between geographically separated clusters. The software supports the following data replication products:

- Sun StorEdge™ Availability Suite
- Hitachi TrueCopy
- EMC Symmetrix Remote Data Facility

The Sun Cluster Geographic Edition software supports Oracle Real Application Clusters with Hitachi TrueCopy software.

The Sun Cluster Geographic Edition product provides highly available services within a cluster by utilizing Sun Cluster resource management features.

Administration and Configuration Tools

You can configure, control, and monitor partnerships, heartbeats, and protection groups either through the SunPlex Manager Geographic Edition GUI or through the command-line interface (CLI).

The Sun Cluster Geographic Edition CLI contains a set of dedicated commands.

The SunPlex™ Manager Geographic Edition GUI is an extension of the Sun Cluster GUI. The GUI provides a visual display for most of the operations that are available through the CLI. You can manage and monitor geographically separated clusters by using the SunPlex Manager Geographic Edition GUI.

Key Concepts for Sun Cluster Geographic Edition

This chapter describes the key concepts for using the Sun Cluster Geographic Edition product. These concepts can help you understand the relationships between the Sun Cluster Geographic Edition components.

This chapter contains the following sections:

- “Data Replication” on page 13
- “Cluster Partnerships” on page 15
- “Protection Groups” on page 16
- “Heartbeat Monitoring” on page 18

Data Replication

Data replication enables controlled migration of production services from a primary cluster to a secondary cluster either in the event of a disaster or as part of a planned procedure. Data is continuously replicated from the primary cluster to the secondary cluster either synchronously or asynchronously, or a combination of both, depending on the recover point objectives of the application services that are supported by the clusters.

The Sun Cluster Geographic Edition software supports Sun StorEdge Availability Suite 3.2.1 software, Hitachi TrueCopy, and EMC Symmetrix Remote Data Facility software for data replication. The Sun StorEdge Availability Suite 3.2.1 software uses a host-based data replication facility which replicates data at the file system or logical volume level within the operating system. The Hitachi TrueCopy software and EMC Symmetrix Remote Data Facility software use a storage-based data replication facility which replicates data at the storage system level and provides a transparent service to applications.

The Sun StorEdge Availability Suite 3.2.1 software is a host-based replication facility that replicates disk volumes between geographically separated primary clusters and secondary clusters in real time. Remote mirror replication enables data from the master volume of the primary cluster to be replicated to the master volume of the geographically separated secondary cluster through a TCP/IP

connection. A remote mirror bitmap tracks differences between the master volume on the primary disk and the master volume on the secondary disk.

The remote mirror software continually replicates the data to remote sites while your applications are accessing the data volumes. You can also manually update the data on the secondary site volume by issuing a command to synchronize the primary and secondary site volumes. You can also restore data from the secondary volume to the primary volume by issuing a command to reverse-synchronize the volumes. For more information about the Sun StorEdge Availability Suite 3.2.1 software, see the product documentation.

The Hitachi TrueCopy software is a storage-based replication facility that provides a host-independent data replication for geographically separated clusters. Hitachi TrueCopy software enables the primary volumes to remain online for all hosts and while processing both read and write I/O operations. If a disaster or system failure occurs, the secondary copy of the data can be run for a recovery with minimal loss of data. For more information about the Hitachi TrueCopy software, see the product documentation.

The EMC Symmetrix Remote Data Facility software provides remote storage replication for disaster recovery and ensures data and systems availability through remote site failover. EMC Symmetrix Remote Data Facility devices are configured in pairs. The mirroring relationship between the pairs becomes operational as soon as the EMC Symmetrix Remote Data Facility links are online. The EMC Symmetrix Remote Data Facility global memory stores information about the pair state of operating EMC Symmetrix Remote Data Facility devices. For information about the EMC Symmetrix Remote Data Facility software, see the product documentation.

Resource groups and device groups enable the Sun Cluster Geographic Edition software to manage data replication and takeover between clusters. You can also configure a protection group to replicate data from a primary cluster to a secondary cluster. For more information about configuring data replication, see the guide for the data replication product you are using:

- Chapter 1, “Replicating Data With Sun StorEdge Availability Suite 3.2.1 Software,” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorEdge Availability Suite*
- Chapter 1, “Replicating Data With Hitachi TrueCopy Software,” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- Chapter 1, “Replicating Data With EMC Symmetrix Remote Data Facility Software,” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

Replication Resource Groups

The Sun Cluster Geographic Edition software extends Sun Cluster resource management features to integrate the data replication products. When you configure a protection group, the Sun Cluster Geographic Edition software creates replication resource groups for monitoring and controlling data replication.

Device Groups

A device group is a hardware resource that Sun Cluster manages. A device group is a type of global device that the Sun Cluster software uses to register volume manager disk groups. The Sun Cluster Geographic Edition software configures Sun Cluster device groups to include replication. For information about configuring device groups in Sun Cluster, see the guide for the data replication product you are using:

- Chapter 1, “Replicating Data With Sun StorEdge Availability Suite 3.2.1 Software,” in *Sun Cluster Geographic Edition Data Replication Guide for Sun StorEdge Availability Suite*
- Chapter 1, “Replicating Data With Hitachi TrueCopy Software,” in *Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy*
- Chapter 1, “Replicating Data With EMC Symmetrix Remote Data Facility Software,” in *Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility*

Cluster Partnerships

A partnership establishes heartbeat monitoring between two clusters that are running Sun Cluster Geographic Edition software. Clusters in a partnership exchange heartbeats to monitor each other’s presence and health. You can configure a partnership between only two clusters, and only one partnership can be defined between the clusters. The two clusters must have an Internet connection with each other. A partnership establishes heartbeats between the clusters.

The Sun Cluster Geographic Edition software uses the IP interconnect between partner clusters for management as well as heartbeats. The IP interconnect can be secured via IPsec for additional security when a public network is in use.

The Sun Cluster Geographic Edition software enables you to specify a command to execute when a heartbeat-loss notification is issued. This command is executed with root permissions. You can also specify a list of email addresses to notify when a heartbeat-loss notification has been issued.

The following figure illustrates a partnership between two clusters.




FIGURE 2-1 Partnerships Between Clusters

A single cluster can participate in more than one partnership with other clusters, but two clusters cannot be in more than one partnership with each other.

Protection Groups

Protection groups enable a set of clusters to tolerate and recover from disaster by managing the resource groups for services. Protection groups can exist only in a partnership. You must create a partnership before you can create a protection group for that partnership. One partner cluster is the primary cluster of the protection group, and the other partner cluster is the secondary cluster. A protection group contains application resource groups and properties for managing data replication for those application resource groups. You must duplicate the application resource group configuration on partner clusters. The configuration for a protection group is identical on partner clusters, so partner clusters must have the application resource groups of the protection group defined in their configuration. The Sun Cluster Geographic Edition software propagates protection group configurations between partners.

You can specify a data replication type in the protection group to indicate the mechanism that is used for data replication between partner clusters. A protection group supports only one data replication type. A protection group can manage one or more application resource groups. When a service is protected from disaster by data replication, the protection group also contains replication resource groups. Protection groups link an application in a resource group with the application data that should be replicated. This linkage and replication enable the application to fail over seamlessly from one cluster to another cluster.

Relationship Between Partnerships and Protection Groups

Clusters in a protection group must be defined as partners. Protection groups require a partnership that defines the clusters that can host the protection group. A cluster can be defined in more than one protection group, and the cluster can have a different role in each protection group. For example, the primary cluster of one protection group can also be the secondary cluster of another protection group. A partnership can have any number of protection groups.

The following figure illustrates two clusters that are defined in one cluster partnership and two protection groups.

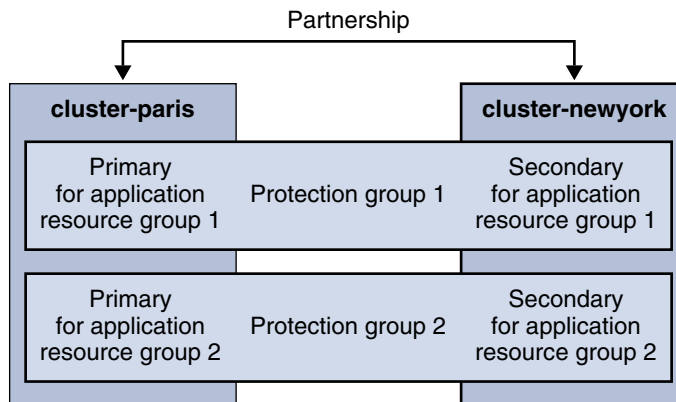


FIGURE 2-2 Example Configuration of Two Clusters in Protection Groups

The following figure illustrates three clusters that are defined in two cluster partnerships and two protection groups.

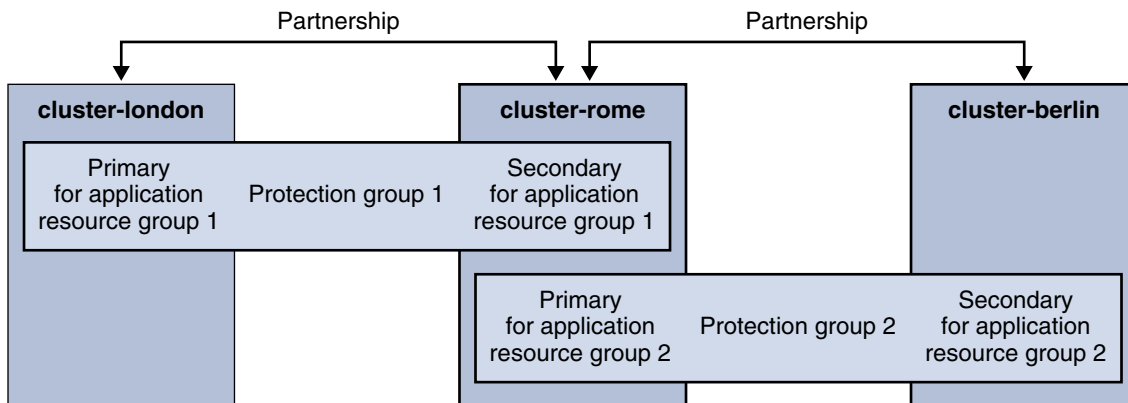


FIGURE 2-3 Example Configuration of Three Clusters in Protection Groups

Protection Group Status

The Sun Cluster Geographic Edition software monitors the status of a protection group on each cluster. The software then combines the local status of each cluster into a global view of the protection group status. The global status reflects the overall status of the protection group.

You can view the protection group status from the SunPlex Manager GUI or through the CLI.

For more information about the status of protection groups, see the *Sun Cluster Geographic Edition System Administration Guide*.

Application Resource Groups

To be highly available, an application must be managed as a resource in an application resource group. You can configure an application resource group for a takeover application or a scalable application. You must also configure application resources and application resource groups on both the primary cluster and the secondary cluster. The data that the application resource accesses must be replicated on the secondary cluster.

The replication for the data volumes that an application resource accesses must be in the same protection group as the applications.

Support for data replication might limit how you configure application resource groups. These requirements and limitations vary with the data replication type you choose. For more information about these requirements, see the *Sun Cluster Geographic Edition System Administration Guide*.

Heartbeat Monitoring

The Sun Cluster Geographic Edition software uses heartbeats to monitor the state between partner clusters. Heartbeats are sent over the public network to detect cluster failures at geographically separated sites. Heartbeat monitoring is part of a partnership configuration. For example, a cluster failure occurs when all nodes of a cluster shut down. The Sun Cluster Geographic Edition software uses the heartbeat status to notify administrators of failures or to trigger a failover to a secondary cluster at an alternate site. Heartbeats could also be lost when a cluster loses access to the public network and no communication occurs between the partner clusters.

Heartbeat Plug-ins

The heartbeat monitor uses plug-in modules to query the heartbeat status of its partners. The Sun Cluster Geographic Edition software offers default plug-ins for monitoring through a TCP/UDP connection.

You can use customize plug-ins to provide a data path over alternate communication links, such as email, HTTP, satellite, and microwave towers.

Sun Cluster Geographic Edition Architecture

The Sun Cluster Geographic Edition software enables a group of clusters to be managed and viewed as a single, large system. This chapter presents a high-level overview of the Sun Cluster Geographic Edition architecture, which you can use in preparation for installing, configuring, and administering Sun Cluster Geographic Edition software.

This chapter contains the following topics:

- “Sun Cluster Geographic Edition Software Environment” on page 19
- “Sun Cluster Geographic Edition Hardware Environment” on page 21
- “Data Replication Configuration” on page 21
- “Geographically Distributed Cluster Topology” on page 23

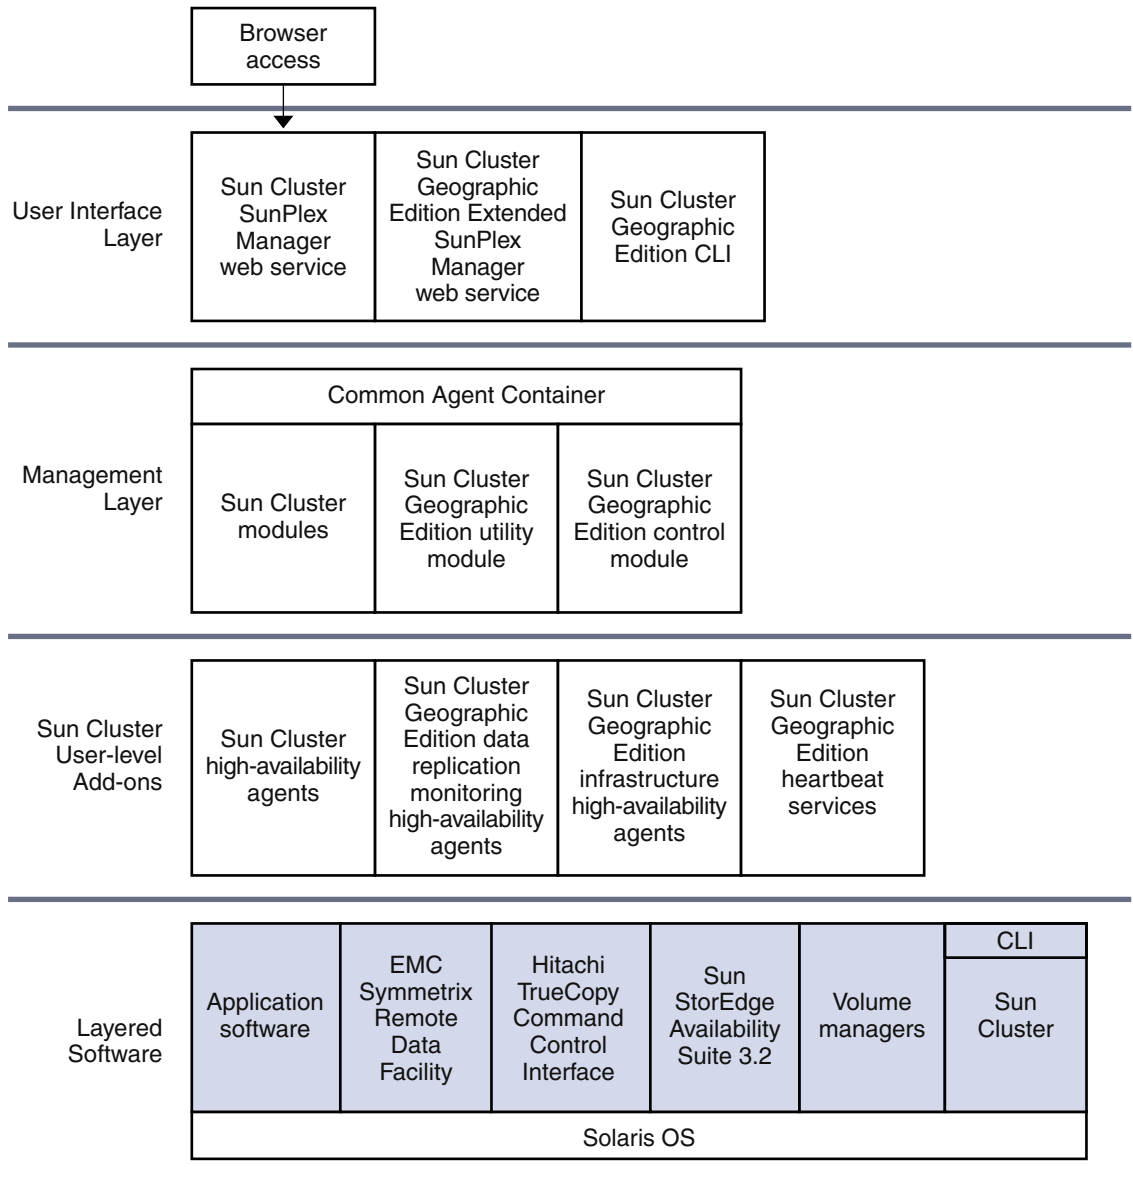
Sun Cluster Geographic Edition Software Environment

The Sun Cluster Geographic Edition software provides tools for managing geographically separated clusters. The Sun Cluster Geographic Edition product also provides highly available services within a cluster by utilizing Sun Cluster resource management features.

The following software components form a Sun Cluster Geographic Edition cluster:

- Solaris™ 8, 9, or 10 software
- Sun Cluster software
- Sun Cluster Geographic Edition software
- Application Data Service Agents
- Data replication software
- Either Solaris Volume Manager or Veritas Volume Manager

The following figure illustrates how components of the Sun Cluster Geographic Edition software interact.



 Platform Software
 Figure 3-1 Overview of Sun Cluster Geographic Edition Software Architecture

The Sun Cluster Geographic Edition software can be installed and removed independent of the underlying Sun Cluster installation. The installation and the uninstallation processes do not require an additional node reboot or cluster downtime.

Sun Cluster Geographic Edition Hardware Environment

The Sun Cluster hardware configuration is the basis for a Sun Cluster Geographic Edition cluster.

The following additional hardware components form a Sun Cluster Geographic Edition cluster:

- Sun Cluster hardware installations, with attached data storage
- Internet connections for inter-cluster management communication between the Sun Cluster installations
- Internet connections for inter-cluster heartbeats
- Connections for data replication
- Connections for custom heartbeats

The Sun Cluster Geographic Edition hardware environment supports the following topologies:

- N+1, where multiple clusters that are located at multiple sites are communicating with a single backup cluster
- Cluster pair, where both clusters are online and providing services

Figure 3–3 illustrates a high-level view of the Sun Cluster Geographic Edition hardware architecture.

Data Replication Configuration

The Sun Cluster Geographic Edition software does not limit the distance between partner clusters. Partner clusters require data replication connections to support the protection groups that are hosted by the partnership. Partner clusters must be compatibly configured to support data replication between the clusters.

The Sun Cluster Geographic Edition software supports replication from a single-node cluster to a single-node cluster, from a multinode cluster to a single-node cluster, and from a multinode cluster to a multinode cluster.

While you can use a single-node cluster at both the primary and backup sites, a single-node cluster offers no internal redundancy. To ensure no single point of failure, you must have a minimum of two nodes in a cluster at the primary site. Use a single-node cluster at the secondary site as a cost effective backup solution if the secondary site is used only for backup purposes. Do not use single-node cluster to run mission critical applications.

Primary clusters and secondary clusters can have any configuration that is supported by the Sun Cluster product if the data replication characteristics of these clusters are compatible. The level of compatibility varies with each data replication product.

The following requirements determine the data replication connection:

- The distance between the partner clusters
- The amount of data that is being replicated

- Data replication configuration parameters

The Sun Cluster Geographic Edition product enables you to balance between data consistency and the cost of the connection, where data consistency is the acceptable amount of data loss.

The following figure illustrates a data replication configuration from a two-node cluster to a single-node cluster.

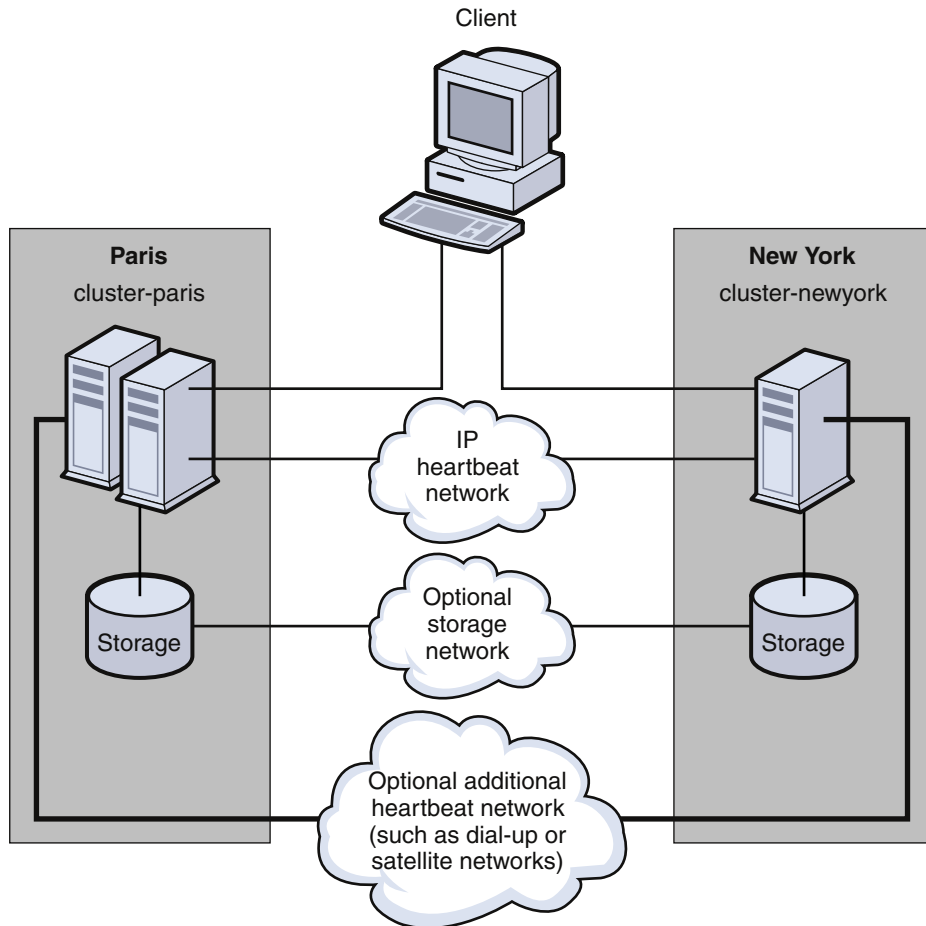


FIGURE 3-2 Data Replication From a Two-Node Cluster to a Single-Node Cluster

The following figure illustrates a data replication configuration from a two-node cluster to a two-node cluster.




FIGURE 3-3 Data Replication From a Two-Node Cluster to a Two-Node Cluster

Geographically Distributed Cluster Topology

A partnership establishes communication and a heartbeat between clusters. One cluster can participate in several partnerships. A protection group establishes data replication between partner clusters. You can configure several protection groups for a partnership, with each protection group replicating different data between the partner clusters.

The following figure illustrates a geographically distributed topology that demonstrates intercluster relationships.





FIGURE 3-4 Geographically Distributed Topology

The Sun Cluster Geographic Edition software enables you to configure multiple partnerships for a cluster with heartbeats between partner clusters. For example, the Geneva-London-Rome-Berlin topology contains a central cluster in Geneva that forms three separate partnerships with clusters in London, Rome, and Berlin. The partnerships require two-way Internet connections between the following cluster pairs: London and Geneva, Rome and Geneva, and Berlin and Geneva. These partnerships enable the cluster in Geneva to detect failures of the clusters in London, Berlin, and Rome by exchanging heartbeats.

Each partnership has a protection group so that the primary clusters in London, Rome, and Berlin can replicate data to the cluster in Geneva as a secondary cluster.

The following figure illustrates a geographically distributed topology that demonstrates intercluster relationships.



The Paris-New York topology has two clusters that form a partnership with two protection groups. Each cluster is the primary cluster for one protection group, and the secondary cluster for the other protection group. The partnership requires a two-way Internet connection between the two clusters for intercluster management and heartbeats. The two clusters must have a data replication link to support data replication for two protection groups.

In the Geneva-London-Rome-Berlin topology, the cluster in Geneva could become the primary cluster for any of the three protection groups. However, the cluster in Geneva must have adequate provisioning to run all the services that are offered by the application resource groups.

For example, if the cluster in Rome is shut down for maintenance, the cluster in Geneva could be the new primary cluster by using a controlled switchover for the Rome-Geneva protection group. As the new primary cluster for the Rome-Geneva protection group, the cluster in Geneva would host the services that are provided by the application resource groups of the Rome-Geneva protection group. The cluster in Geneva would simultaneously serve as the secondary cluster for the clusters in London and Berlin.

Similarly, in the Paris-New York topology, either cluster could take over as the primary cluster to both protection groups if the partner cluster were unexpectedly lost.

Glossary

active–active clusters	Two clusters where each cluster is both the primary cluster for some services and the secondary cluster for other services.
application resource	An application that is managed as a resource to its increased availability.
application resource group	A Sun Cluster resource group that is configured by the user to make an application highly available on Sun Cluster. An application resource group can be configured into a protection group to make it disaster tolerant and highly available.
campus cluster	A cluster that supports data replication between geographically separate nodes within <i>one</i> cluster. The maximum distance between nodes is limited.
data replication	<p>The copying of data from data storage devices in a primary cluster to data storage devices in a secondary cluster. Through data replication, the secondary cluster has a recent copy of the data on the primary cluster. The primary and secondary clusters can be geographically separated.</p> <p>In a campus cluster, the two data storage devices are on the same cluster. In a geographically separated cluster that runs the Sun Cluster Geographic Edition software, the two data storage devices are on different clusters.</p>
data replication resource	A Sun Cluster resource that monitors the state and status of data replication.
disaster tolerance	The ability of a system to restore an application on a secondary cluster when the primary cluster fails. Disaster tolerance is based on data replication and failover.
disconnected partnership	An error scenario in which two clusters in a protection group act as the primary cluster. In a disconnected partnership, the system administrator must execute a takeover, making one cluster the primary cluster and the other cluster the secondary cluster.
HASStoragePlus resource	A resource that enforces the switchover of the device group when the replication resource group is switched over or failed over.
heartbeat	A signal that is emitted from a cluster and detected by its partner cluster. Heartbeats enable a cluster to monitor the presence and failure of its partner cluster.
inactive cluster	A primary cluster where an application is not running and data is not being replicated to a secondary cluster. Alternatively, a secondary cluster where data is not being replicated from the primary cluster.

partnership	A relationship between two geographically separated clusters that are installed with the Sun Cluster software and Sun Cluster Geographic Edition software. These two clusters monitor each other's health by exchanging heartbeats.
primary cluster	A cluster that is in a cluster partnership, that hosts the application resources, and that holds the primary copy of replicated data. Protection groups define whether a cluster is primary or secondary. For example, the primary cluster of one protection group can also be the secondary cluster for another protection group.
protection group	<p>An entity that manages application resource groups for services that are protected from disaster. Clusters in a protection group must be defined as partners. A cluster can have different roles in different protection groups. For example, the primary cluster of one protection group can also be the secondary cluster for another protection group.</p> <p>A protection group has the following characteristics:</p> <ul style="list-style-type: none">▪ A set of resource groups and resources for services that are protected from disaster▪ Device group entries▪ A primary cluster that hosts the protection group▪ A secondary cluster that is able to host the protection group▪ A data replication service
replication resource group	A resource group that contains data replication resources.
resource	A Sun Cluster resource.
resource group	A Sun Cluster resource group. A resource group can be an application resource group or a replication resource group.
scalable application	An application that runs on several nodes of one cluster to create a single, logical service. If a node that is running a scalable application fails, failover does not occur. The application continues to run on the other nodes of the cluster.
secondary cluster	<p>A cluster in a cluster partnership that is capable of hosting a protection group. The secondary cluster receives mirrored data from the primary cluster. If the primary cluster fails, the secondary cluster can become the new primary cluster.</p> <p>A secondary cluster can be associated with a protection group. If a primary cluster fails, the protection group is migrated to a secondary cluster. Protection groups define whether a cluster is primary or secondary. For example, the primary cluster of one protection group can also be the secondary cluster for another protection group.</p>
secondary node	A node that is in a cluster, but does not host the application services. If the primary node fails, the secondary node becomes the new primary node.

site	<p>A location that houses one or more clusters that run the Sun Cluster Geographic Edition software. To participate in the disaster recovery environment, a cluster must have a partner cluster on a geographically separated site.</p>
standby cluster	<p>A cluster with a minimal configuration that acts as a secondary cluster. A standby cluster can take over from the primary cluster in an emergency situation, but supports a reduced service only. A standby cluster is a low-cost alternative to a secondary cluster.</p>
switchover	<p>The planned migration of services from the primary cluster to the secondary cluster.</p> <p>Unlike a takeover, the primary cluster is connected to the secondary cluster during a switchover. During a switchover, the primary cluster is connected to the secondary cluster and coordinates the migration of services with the secondary cluster. This coordination enables the data replication to be completed and ensures that services can be transferred from the primary cluster to the secondary cluster with minimal loss or corruption of data.</p>
takeover	<p>The emergency migration of services from the primary cluster to the secondary cluster. A system administrator can initiate a takeover to recover from a disaster scenario.</p> <p>Unlike a switchover, the primary cluster is not connected to the secondary cluster during a takeover. Therefore, the primary cluster cannot coordinate with the secondary cluster to migrate the services. Because of this lack of coordination, the risk of data loss and data corruption is higher than with a switchover. During a takeover, dedicated recovery procedures are used to minimize data loss and data corruption.</p>

Index

A

- applications
 - fault-tolerant, 10
 - resource groups, 18
- architecture
 - hardware, 21
 - software, 19-20

C

- clusters, topologies, 23-25

D

- data replication
 - configuration, 21-23
 - description, 10
 - resource groups, 13-15
 - supported products, 11-12
- data service agents, 19-20
- device groups, 15
- disaster tolerance, 10, 16-18

E

- EMC Symmetrix Remote Data Facility, 11-12

F

- failure, hardware and software, 10-11

H

- hardware
 - architecture, 21
 - failure, 10-11
- heartbeats
 - description, 11-12, 18
 - plug-ins, 18
- Hitachi TrueCopy, 11-12

I

- IPsec, 11-12

M

- monitoring, failure, 10-11

P

- partnerships
 - data replication in, 21-23
 - description, 15-16
 - protection groups and, 16-17
- plug-ins, heartbeats, 18

protection groups
 description, 16-18
 partnerships and, 16-17
 states, 17

R

RBAC, 11-12
recovery, disaster, 10-11
replication resource groups, 14
resource groups
 application, 18
 replication, 14

S

software
 architecture, 19-20
 configuration, 21-23
 failure, 10-11
Solaris Volume Manager, 19-20
SSL, 11-12
states, protection groups, 17
Sun Cluster, 10, 19-20, 21
Sun StorEdge Availability Suite 3.2.1 software, 11-12,
 19-20

T

topologies, 23-25

V

Veritas Volume Manager, 19-20