HA Cluster Glossary

THE NETWORK IS THE COMPUTER™

SunSoft, Inc.
A Sun Microsystems, Inc. Business
2550 Garcia Avenue
Mountain View, CA 94043 USA
415 960-1300 fax 415 969-9131

Part No.: 802-7073-10
Revision A, November 1996

Please
Recycle

Contents

Glossary

Asymmetric configuration

A configuration that contains a single diskset. In an asymmetric configuration, one server acts as the default master of the diskset and the other server acts as a hot standby.

Cluster reconfiguration

An ordered multistep process that is invoked whenever there is a significant change in cluster state, such as takeover, switchover, or a physical host reboot. During cluster reconfiguration, the Solstice HA software coordinates all of the physical hosts that are up and communicating. Those hosts agree on which logical host(s) should be mastered by which physical hosts.

Concatenation

A metadevice created by sequentially mapping blocks on several physical slices (partitions) to a logical device. Two or more physical components can be concatenated. The slices are accessed sequentially rather than interlaced (as with stripes).

Data service

A network service that implements read-write access to disk-based data from clients on a network. NFS is an example of a data service. The data service may be composed of multiple processes that work together.

Default master

The server that is master of a diskset if both servers rebooted simultaneously. The default master is specified when the system is initially configured.

Diskset

A group of disks that move as a unit between the two servers in a HA configuration.

Fault detection

Solstice HA programs that detect two types of failures. The first type includes low-level failures such as system panics and hardware faults (that is, failures that cause the entire server to be inoperable). These failures can be detected quickly. The second type of failures are related to data service, such as HA-NFS. These types of failures take longer to detect.

HA Administrative file system

A special file system created on each logical host when Solstice HA is first installed. It is used by Solstice HA and by layered data services to store copies of their administrative data.

HA-NFS

Highly available NFS (Sun's distributed computing file system). HA-NFS provides highly available remote mount service, status monitor service, and network locking service.

Heartbeat

A periodic message sent between the two membership monitors to each other. Lack of a heartbeat after a specified interval and number of retries may trigger a takeover.

Highly available data service

A data service that appears to remain continuously available, despite single-point failures of server hardware or software components.

Hot standby

In an asymmetric (single diskset) configuration, the machine that is not the current master of the diskset. If both servers reboot simultaneously, the hot standby will not master the diskset and thus will not be running any Solstice HA data services.

Local disks

Disks attached to a HA server but not included in a diskset. The local disks contain the Solaris distribution and the Solstice HA and DiskSuite software packages. Local disks must not contain data exported by the Solstice HA data service.

Logical host

A diskset and its collection of logical host names and their associated IP addresses. Each logical host is mastered by one physical host at a time.

Logical host name

The name assigned to one of the logical network interfaces. A logical host name is used by clients on the network to refer to the location of data and data services. The logical host name is the name for a path to the logical host. Because a host may be on multiple networks, there may be multiple logical host names for a single logical host.

Logical network interface

In the Internet architecture, a host may have one or more IP addresses. HA configures up additional logical network interfaces to establish a mapping between several logical network interfaces and a single physical network interface. This allows a single physical network interface to respond to multiple logical network interfaces. This also enables the IP address to move from one HA server to the other in the event of a takeover or `haswitch`(1M), without requiring additional hardware interfaces.

Master

The server with exclusive read and write access to a diskset. The current master host for the diskset runs the data service and has the logical IP addresses mapped to its Ethernet address.

Membership monitor

A process running on both HA servers that monitors the servers. The membership monitor sends and receives heartbeats to its sibling host. The monitor can initiate a takeover if the heartbeat stops. It also keeps track of which servers are active.

Metadevice

A group of components accessed as a single logical device by concatenating, striping, mirroring, or logging the physical devices. Metadevices are sometimes called pseudo devices.

Metadevice state database

Information kept in nonvolatile storage (on disk) for preserving the state and configuration of metadevices.

Metadevice state database replica

A copy of the state database. Keeping multiple copies of the state database protects against the loss of state and configuration information. This information is critical to all metadevice operations.

Mirroring

Replicating all writes made to a single logical device (the mirror) to multiple devices (the submirrors), while distributing read operations. This provides data redundancy in the event of a failure.

Multihomed host

A host that is on more than one public network.

Multihost disk

A disk configured for potential accessibility from multiple servers. Solstice HA software enables data on a multihost disk to be exported to network clients via a highly available data service.

Sibling host

One of the two physical servers in a HA configuration.

Solstice HA

See Solstice High Availability.

Solstice High Availability

A software package that enables two machines to act as read-write data servers while acting as backups for each other.

Solstice DiskSuite

A software product that provides data reliability through disk striping, concatenation, mirroring, UFS logging, dynamic growth of metadevices and file systems, and metadevice state database replicas.

SPARCcluster High Availability

The combination of Solstice HA software with Sun hardware creating a highly available system.

Stripe

Similar to concatenation, except the addressing of the component blocks is non-overlapped and interlaced on the slices (partitions), rather than placed sequentially. Striping is used to gain performance. By striping data across disks on separate controllers, multiple controllers can access data simultaneously.

Submirror

A metadevice that is part of a mirror. See also mirroring.

Switchover

The coordinated moving of a logical host (diskset) from one operational HA server to the other. A switchover is initiated by an administrator using the `haswitch`(1M) command.

Symmetric configuration

A HA configuration that contains two disksets. In a symmetric configuration, each server is the default master for one diskset.

Takeover

The automatic moving of a logical host from one HA server to the other after a failure has been detected. The HA server that has the failure is forced to give up mastery of the logical host.

Trans device

A pseudo device responsible for managing the contents of a UFS log.

UFS

An acronym for the UNIX® file system.

UFS logging

Recording UFS updates to a log (the logging device) before the updates are applied to the UFS (the master device).

UFS logging device

The component of a transdevice that contains the UFS log.

UFS master device

The component of a transdevice that contains the UFS file system.