Datasheet

Linux FailSafe™ 1.0

Features

- The most advanced, full-featured high-availability solution for Linux at a low cost
- Open-source software and API for ISV application integration
- Support for SuSE and Red Hat
- Multinode, all active, N+1 or N+M clusters
- Easy administration with Java language–based GUI
- Dynamic cluster reconfiguration, fine-grained and cascading failover for minimal disruption to highly available services
- Recovery from network partitioning
- Support for XFS, ext3, and other journaled filesystems
- Planned integration with other open-source Linux clustering facilities such as Linux Virtual Server and network mirroring

LINUX

The High-Availability (HA) Solution for Business-Critical Applications

SGI™ IRIS FailSafe™, the high-availability solution for SGI IRIX®, is now available for Linux® as Linux FailSafe. IRIS FailSafe 2.0 has been extensively deployed across industries such as telecommunications, entertainment, manufacturing, and government for services that must be highly available, including file serving, database serving, HSM, and other custom applications. Linux FailSafe now brings the same proven enterprise-class HA solution to our Linux customers.

In contrast to Internet load balancing HA solutions for Web farms, Linux FailSafe brings you an advanced, reliable, high-availability solution for business-critical applications that must be available around the clock. With Linux FailSafe, critical applications and services can now be deployed and run without interruption on Linux platforms, much like they do on the robust UNIX® platforms.

Linux FailSafe Architecture

In a Linux FailSafe cluster, multiple servers are connected to both a public and a private network. Clients use the public network to access services from this cluster of Linux systems. The Linux FailSafe software uses the private network to monitor the cluster members. Following are some of the salient features of Linux FailSafe:

- All systems can be active in normal operation; no redundant investment is necessary
- In case of a system or application failure, user-configured failover policies determine whether the application will be restarted on the current system or whether one of the other systems will take over the workload
- Fine-grained resource group–based failover minimizes the impact of any failure to only those services directly affected by the failure
- Cascading failovers allow additional protection from multiple system failures
- Efficient run-time addition and deletion of systems and applications minimize the need for planned system downtime for cluster reconfiguration/upgrades
- Quorum facilities help recover from network partitioning

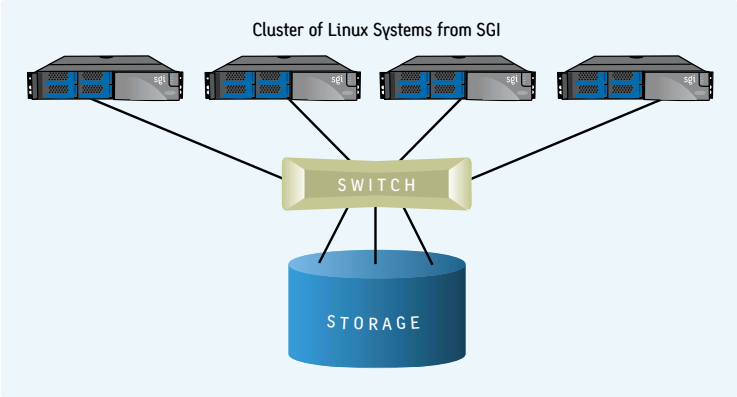

The clustered nodes share storage on either RAIDs or mirrored disks. In the event of a failure, a shared storage subsystem allows multiple servers to assume control of the data; the filesystems are automatically made available on the system(s) where their corresponding applications are resumed.

High Availability at a Low Cost

Linux FailSafe brings you high availability at a fraction of the cost of specialized fault-tolerant systems. It uses standard off-the-shelf components such as Ethernet hubs for interconnecting multiple Linux OS–based systems in a cluster, and there is no redundant investment in a Linux FailSafe cluster.

All systems in the cluster are active in normal operation, working as if they were independent servers. In the event of a server failure, any of the other systems can take over the services of the failed system, transparently fulfilling requests from clients on the network.

Easy Setup and Administration

A Linux FailSafe cluster is configured and managed using the intuitive Java™ language–based Cluster Manager GUI. It has the same features and look and feel as the IRIS FailSafe Cluster Manager. The GUI can be run from any Java language–compliant browser, giving you the utmost flexibility in choosing your cluster management platform. Ported directly from the IRIS FailSafe Cluster Manager, it brings all its capabilities, such as dynamic cluster configuration and status updates, for easy cluster management. A text-based command line interface provides administration over slow connections and script-based automation.

Meets the Needs of Critical Business Services

Linux FailSafe is designed to provide common system and cluster-level resource monitoring and recovery services such as I/O subsystems and IP address failover to all applications. A simple API is provided to develop application agents for integrating any crash-tolerant application with the Linux FailSafe infrastructure to minimize application downtime due to system or application failures.

The Open Source Advantage

SGI has made this advanced technology available to the open source community and is actively encouraging the participation of the wider community for support and continuing enhancements to Linux FailSafe. The programming API has been made available for ISVs to integrate their applications into the Linux FailSafe HA framework. This clustering technology has the potential to provide a strong foundation for building advanced scalable HA cluster solutions for Linux platforms.

World-Class SGI Service

Having been extensively involved in custom application integration with IRIS FailSafe, the highly qualified SGI Professional Services team is available to develop agents for integrating custom or third-party applications into the Linux FailSafe HA framework, extending the high-availability functionality to a wide range of applications.

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America 1(650) 933-4637
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 771.0290

©2000 Silicon Graphics, Inc. All rights reserved. Specifications subject to change without notice. Silicon Graphics, IRIX, and IRIS are registered trademarks, and SGI, Linux FailSafe, XFS, IRIS FailSafe, and the SGI logo are trademarks, of Silicon Graphics, Inc. Red Hat is a registered trademark of Red Hat, Inc. Java is a trademark of Sun Microsystems, Inc. Linux is a registered trademark of Linus Torvalds. UNIX is a registered trademark in the U.S. and other countries, licensed exclusively through X/Open Company Limited. All other trademarks mentioned herein are the property of their respective owners.

2872 [8/2000]

J11699