

SunPlex™ Systems and Sun™ Cluster 3.1 Software

A premier application availability platform.



Key feature highlights

Integrated With the Solaris™ Operating System: Key fault detection and recovery functions of the SunPlex™ system are integrated with the Solaris OS — leading to fast failure detection and recovery.

High Availability and Reduced Risk: IT can deliver higher, more predictable application service levels while reducing IT risks and costs. The SunPlex system automatically restarts a failed resource without the need for manual intervention, enabling consistent application service levels.

Security: The SunPlex environment is a security hardened environment, so application services on a SunPlex environment are safer from hacker intrusions and denial of service attacks.

Great Flexibility in Configurations: Sun Cluster 3 offers superior flexibility by supporting server-storage connections in any configuration, and allows the mixing of any combination of servers and CPUs in the same cluster as well as the mixing of UltraSPARC® II and UltraSPARC III systems in the same cluster.

Ease of Use: Cluster administrative and development tasks can be performed with ease; administrators can manage any resource on a SunPlex system from anywhere on the network and application developers can cluster-enable applications in two clicks with the SunPlex Agent Builder.

Disaster Tolerance: Using Sun's Enterprise Continuity solutions, nodes in the SunPlex environments can be separated by up to 200 Km for continuity of services in the event of a catastrophic failure.

Rolling Upgrade Framework: Sun Cluster 3.1 contains a software versioning framework, that will allow upgrades from Sun Cluster 3.1 to later releases without having to shut down the entire cluster.

Businesses today demand application service continuity at a low cost. The SunPlex™ system offers a unique approach to application service-level management, and is the industry's premier availability platform for improving the predictability and resilience that businesses expect from their application infrastructure.

Sun™ Cluster 3 software powers SunPlex systems to manage application services, optimizing both the availability and scalability of these services. Sun Cluster 3 is the key to building highly available SunPlex systems that include the Solaris™ Operating System (OS), agents and applications, Sun servers, Sun StorEdge™ solutions, interconnect products, and services. Since its introduction in November 2000, Sun Cluster 3 software and follow-on update releases have been widely adopted in the market. With this latest release, Sun Cluster 3.1 contains features that further enhance the availability and manageability characteristics of a customer's infrastructure.

Sun Cluster 3.1 and the Solaris OS

Sun Cluster 3 software takes general-purpose clustering beyond the realm of high availability by adding the simplicity of single-system manageability and the potential of seamless scalability. In essence, the cluster becomes a single managed entity, and presents itself and its services to clients as if it were an individual server.

The Sun Cluster 3 framework extends the Solaris OS, enabling core Solaris services — devices, file systems, and networks — to operate seamlessly across a SunPlex system while still maintaining full Solaris OS compatibility with existing applications. Sun Cluster 3 provides high availability (HA) and scalability to everyday Solaris applications through continuous network and data availability. Applications that have agents written to the easy-to-use Sun Cluster 3 API can achieve even higher levels of availability as well as scalability.

SunPlex System Key Benefits

The SunPlex system offers a unique, cost-effective approach to application service-level management, and is the premier, end-to-end, integrated application service delivery platform. With the help of SunPlex systems, enterprises can increase application service levels while decreasing service-level cost and risk.

Application Service Levels Easier to Manage
SunPlex systems provide a powerful, flexible environment where multiple application services can be simultaneously deployed and easily managed. Application services running on SunPlex systems receive the full benefits of the Sun Cluster 3 HA framework. Applications can be deployed across the SunPlex system without worrying about how they will find and access network and file services.

The SunPlex system is a premier, end-to-end, integrated application service delivery platform. Combined with Sun Cluster 3 software, SunPlex systems are designed to manage application services for tightly coupled environments, optimizing both the availability and scalability of these services.

Continuous Access to Data and Network

SunPlex systems deliver Global Network Services and Global File Services. Data, network, and devices are continuously available to all domains in the SunPlex system as well as to applications running on any domain. Sun Cluster 3.1 provides support for Solaris IP Multipathing (IPMP), allowing SunPlex deployments to benefit from the added resilience and improved data throughput.

Simplified Administration, Low Costs

SunPlex systems incorporate system management tools — Sun Management Center and SunPlex Manager tool — to create a centrally managed environment that delivers easy administration and lower operating costs. The Sun Cluster module in Sun Management Center enables a third-party SNMP agent to collect configuration and state information about Sun Cluster software, providing the added flexibility of managing the SunPlex environment via third-party management tools.

Summary of Features

Global Network Services

In the Sun Cluster 3 architecture, incoming requests from the network go to a global interface — a network interface card that hosts the global IP address. The requests are then load-balanced to the various instances of the distributed application running within the cluster. Outgoing packets go out to the network through the local network interface card to prevent saturation of the global interface. In the event of a failure, the global IP is failed over to a backup network interface card. In this manner, the SunPlex Global Network Service provides a highly available global IP address, as well as the simplicity of a single system to clients.

Global Devices and Global File Services

Data access is significantly enhanced in Sun Cluster 3 with the addition of Global Devices and Global File Services. With Global Devices, every domain has access to any device on the SunPlex system — such as a disk or CD-ROM drive — even if that device is not physically connected to that domain. Global File Services extend these capabilities by using shared storage devices — storage with physical connections to more than one domain — so that data is both highly available and accessible to application services running on any domain in the SunPlex system.

Centralization of Global File Services on behalf of the SunPlex system facilitates a simple “single-point-of-management” paradigm.

Customers who prefer to fail over the file system can use the Failover File Service feature (available in Sun Cluster 3.0 5/02 and later releases) instead of the Global File Service. Customers also have a choice of UFS and VERITAS VxFS as the file system to be used in the SunPlex environment.


Scalable Services

The Sun Cluster 3 framework allows a distributed application to run within cluster control, providing the ease of manageability as well as automatic recovery of service levels. Instances of a distributed application can be installed and brought online or offline on multiple cluster nodes with a single procedure, saving time and reducing the complexity associated with distributed applications. Distributed applications also have the option of using the SunPlex Global Network Service to receive the benefit of free-of-charge load balancing and a highly available IP.

SunPlex systems provide commonly used load-balancing schemes such as round-robin and sticky. In addition, client affinity is maintained such that transaction requests from a client machine are always sent to the same cluster node. Storing application configuration data on the SunPlex Global File Service enables faster recovery of failed application instances. By adding more domains or systems to the SunPlex system, capacity and continuity are increased. Service levels are maintained in the event of any number of potential outages — planned or unplanned.

Failover Services

The Sun Cluster 3 architecture delivers inherent HA services. It enables IT organizations to maintain service levels on critical applications and services. Failover Services provide high availability to single-instance applications by failing over the application to a backup node in the event of a failure.



Very Fast Recovery from Failure

Sun Cluster 3 provides fast error detection, fast software switchover, and parallelized application and infrastructure restarts. When a failover occurs, the failed instance of the application is automatically restarted on the same node if the node is healthy, or on the backup node. After the failover is complete, the physical server from which the clients of the application get service and data is transparent.

Abstracted Resources

With the foundation of these key abstracted capabilities — Global Network Services, Global Devices, and Global File Services — there is no concept of logical host in Sun Cluster 3 software. Services need not reside on servers physically attached to storage devices, which means service failover does not require storage device or file system failover. Abstracted resources enhance flexibility as well as minimize failover time.

Centralized System Administration and Management Tools

The Sun Cluster 3 framework simplifies administration by enabling SunPlex system resources to be managed and administered as if they were on a single system. Administrators have access to all the nodes of the cluster from anywhere on the network. And because of the integration with the Solaris Operating System, familiar Solaris commands execute just as if only a single system were being administered.

SunPlex system management is accomplished through either a command-line interface (CLI) or GUI-based management tools (Sun Management Center or SunPlex Manager). These GUI-based tools allow complex tasks to be performed with ease by enabling system administrators to manage any resource on a SunPlex system from anywhere on the network. This provides tremendous cost savings for organizations where administrators are responsible for systems located in different buildings, cities, and even countries.

In Sun Cluster 3.1 software, third-party SNMP agents can collect configuration and state information about the cluster by talking to the Sun Management Center agent. This feature provides the added flexibility of using third-party management tools to manage the SunPlex environment.

Easy Agent Development Environment

Sun provides a list of qualified Scalable and HA (Failover) Agents. Alternatively, developers can use the SunPlex Agent Builder to develop Scalable or HA Agents. The SunPlex Agent Builder generates agent code in Ksh or C with two simple clicks. For fast agent deployment with no code writing or modification, developers can also use the Generic Agent functionality of the Agent Builder. The Generic Agent functionality generates a precompiled agent binary that shortens the agent deployment cycle. The Generic Agent functionality also offers tunable parameters that allow for customization of the Generic Agent.

Fast Application Messaging Via Remote Shared Memory Technology

Sun Cluster 3 includes Remote Shared Memory (RSM) technology, which offers improved service levels for distributed applications running in the Sun Cluster environment. The RSM API offered in the Solaris 8 10/01 OS and later releases enables application developers to bypass the TCP/IP stack and access high-speed, high-bandwidth, and low-latency interconnect hardware directly for fast messaging in the Sun Cluster environment. The interconnects that are supported with RSM are Scalable Coherent Interconnect (SCI-PCI) and Sun Fire™ Link Interconnect. RSM technology in Sun Cluster 3 has been tuned for optimal performance of Oracle® RAC.

Dynamic Reconfiguration Support

Sun Cluster 3 software allows for the dynamic additions or removals of hardware resources such as processors, memory, and I/O devices. Dynamic Reconfiguration (DR) support checks

the safety of a DR operation and rejects any unsafe operations.

Security Hardened

Security hardening is supported on all of the Sun Cluster 3 supported agents with the exception of the BroadVision agent. The DB2 agent from IBM is also security hardened.

Intelligent Resource Management Via Prioritized Service Management

Sun Cluster 3 provides Prioritized Service Management (PSM). PSM is a policy-based, service-level management feature that provides high service levels for a high-priority service in the event of its failover to a backup node. It automatically off-loads low-priority services on the backup node to free resources for the high-priority service. The low-priority services can be shut down or failed over to another node.

Disaster Recovery

Using Sun's Enterprise Continuity solutions, Sun Cluster 3 nodes can be separated by up to 200 Km by leveraging dense wave division multiplexer (DWDM) optical technology to provide application service continuity in the event a catastrophic failure brings down an entire site. Applications can be deployed in failover, active-passive, or active-active configurations within the Enterprise Continuity solution.

Campus Clusters that use standard Fibre Channel technology are also supported in the SunPlex environment, allowing for disaster recovery across the distance of 10 Km. Both two-room (quorum device in the same room as the primary node) and three-room (quorum device in its own room) Campus Clusters are supported by SunPlex systems.

Easy Upgrades

Customers can upgrade from any release of Sun Cluster 3.0 directly to Sun Cluster 3.1 without having to reinstall the Sun Cluster software. This minimizes planned service downtime during the Sun Cluster 3 software upgrade. Sun Cluster 3.1 contains a software versioning

SunPlex™ Systems and Sun™ Cluster 3.1 Software

framework that will allow upgrades from Sun Cluster 3.1 to later releases without the need to shut down the entire cluster.

Faster Deployments With Customer Ready Systems (CRS) Program

The SunPlex solution can be delivered through the Customer Ready Systems (CRS) program, Sun's factory integrated solutions that are designed to deliver faster time to value for Sun customers. The CRS program offers unparalleled ease of deployment benefits. The entire hardware and software that constitutes the customer's deployment platform is installed, configured, and tested in the factory before arriving at the customer's site. This enables customers to deploy highly available services on SunPlex environments with ease and increased time to value.

Eight-Node Support

Sun Cluster 3 supports up to eight nodes in a SunPlex system. Customers can mix and match from Sun's extensive server offerings.

SunPlex System Configurations

A SunPlex system comprises two or more SunPlex system-qualified servers (up to eight are supported); storage products; SunPlex system interconnects; public networks running the Solaris 9 or 8 OS.

Supported Server Platforms

The following server platforms are supported; additional platforms will be available in the future:

- Netra™ 120 t 1120/1125, t 1400/1405, T1 AC200/DC200, 20, and 1280 servers
- Sun Enterprise™ 220R, 250, 420R, 450, 3500, 4500, 5500, 6500, and 10000 servers
- Sun Fire V120, V210, V240, 280R, V480, V880, V1280, 3800, 4800/4810, 6800, 12K, and 15K servers

Sun Cluster Software Coexistence

- Solaris 9 Resource Manager and Solaris Resource Manager for the Solaris 8 OS
- Sun Management Center

Sun StorEdge Traffic Manager Supported Storage Products

The following storage products are supported; additional storage products will be available in the future:

- Sun StorEdge S1, MultiPack, D1000, D2, A1000, A3500 (SCSI/FC), A5000, A5100, A5200, T3, SE 3310, 3910/3960, 6120, 6320, 6910/6960, 9910/9960, and 9970/9980 arrays
- Netra st D130, st A1000, and st D1000 arrays

Supported Third-Party Storage

- EMC Symmetrix¹ and Clarion

Supported SAN Switches

- Sun StorEdge Network FC Switch-8 and Switch-16 and Sun StorEdge Network 2Gb FC Switch-64
- Brocade SilkWorm 2400, 2800, 3800, and 12000
- McData 6064

Supported Interconnects²

The following SunPlex system interconnects are supported; additional SunPlex system interconnects will be available in the future:

- On-board 100 Mb/s/sec port, Sun Quad FastEthernet™, Sun-FastEthernet™, and SunSwift™ adapters
- Sun Gigabit Ethernet adapter
- PCI-Scalable Coherent Interface (SCI)
- Sun Fire Link

Supported Public Networks

The following public networks are supported; additional public networks will be available in the future:

- ATM
- On-board 100 Mb/s/sec port, Sun Quad FastEthernet, Sun-FastEthernet, and SunSwift adapters
- Sun Gigabit Ethernet adapter

Supported SunPlex Topologies

Any configuration of servers and storage is supported.

Supported Volume Managers

- Solaris Volume Manager software
- VERITAS Volume Manager (VxVM)
- VERITAS Volume Manager Cluster Functionality

1. Please contact your Sun Sales Representative for more information.
2. Up to six links of interconnects transferring data in parallel are supported in one SunPlex system.

Learn More

Get the inside story on the trends and technologies shaping the future of computing by signing up for the Sun Inner Circle program. You'll receive a monthly newsletter packed with information on the latest innovations, plus access to a wealth of resources. Register today to join the Sun Inner Circle Program at sun.com/joinic.

To receive additional information on Sun software, products, programs, and solutions, visit sun.com/software.

Supported File Systems

- Solaris UFS
- VERITAS File System (VxFS)

Supported SunPlex Agents

The following agents are available now:

- HA DNS
- HA NFS
- HA Oracle
- Oracle Parallel Server (OPS)
- Oraclelegi RAC
- HA Sybase
- Scalable SAP
- HA SAP liveCache
- HA Siebel
- HA Sun ONE Web Server
- HA Sun ONE Application Server
- HA Sun ONE Message Queue
- HA Sun ONE Messaging Server
- HA Sun ONE Directory Server
- HA Sun ONE Calendar Server
- HA Apache Web/Proxy Server
- HA NetBackup
- HA Solstice Backup™ software
- HA Samba
- HA DHCP
- HA IBM WebSphere MQ
- HA IBM WebSphere MQ Integrator
- HA BEA WebLogic Server
- HA Sun StorEdge Availability Suite
- Scalable Sun ONE Web Server
- Scalable Apache Web/Proxy Server
- Scalable BroadVision One-To-One

Agents Available Through Third Parties

- IBM DB2 (EE and EEE)
- HA Informix Dynamic Server
- Sybase ASE

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 800 786-7638 or +1 512 434-1577 Web sun.com



Sun Worldwide Sales Offices: Africa (North, West and Central) +33-13-067-4680, Argentina +5411-4317-5600, Australia +61-2-9844-5000, Austria +43-1-60563-0, Belgium +32-2-704-8000, Brazil +55-11-5187-2100, Canada +905-477-6745, Chile +56-2-3724500, Colombia +571-629-2323, Commonwealth of Independent States +7-502-935-8411, Czech Republic +420-2-3300-9311, Denmark +45 4556 5000, Egypt +202-570-9442, Estonia +372-6-308-900, Finland +358-9-525-561, France +33-134-03-00-00, Germany +49-89-46008-0, Greece +30-1-618-8111, Hungary +36-1-489-8900, Iceland +354-563-3010, India-Bangalore +91-80-2298989/2295454; New Delhi +91-11-6106000; Mumbai +91-22-697-8111, Ireland +353-1-8055-666, Israel +972-9-9710500, Italy +39-02-641511, Japan +81-3-5717-5000, Kazakhstan +7-3272-466774, Korea +822-2193-5114, Latvia +371-750-3700, Lithuania +370-729-8468, Luxembourg +352-49 11 33 1, Malaysia +603-21161888, Mexico +52-5-258-6100, The Netherlands +00-31-33-45-15-000, New Zealand-Auckland +64-9-976-6800; Wellington +64-4-462-0780, Norway +47 23 36 96 00, People's Republic of China-Beijing +86-10-6803-5588; Chengdu +86-28-619-9333; Guangzhou +86-20-8755-5900; Shanghai +86-21-6466-1228; Hong Kong +852-2202-6688, Poland +48-22-8747800, Portugal +351-21-4134000, Russia +7-502-935-8411, Singapore +65-6438-1888, Slovak Republic +421-2-4342-94-85, South Africa +27 11 256-6300, Spain +34-91-596-9900, Sweden +46-8-631-10-00, Switzerland-German 41-1-908-90-00; French 41-22-999-0444, Taiwan +886-2-8732-9933, Thailand +662-344-6888, Turkey +90-212-335-22-00, United Arab Emirates +9714-3366333, United Kingdom +44-1-276-20444, United States +1-800-555-9SUN or +1-650-960-1300, Venezuela +58-2-905-3800

SUN™ © 2003 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, Netra, Solaris, Solstice Backup, Sun Enterprise, SunFastEthernet, Sun Fire, SunPlex, Sun Quad FastEthernet, Sun StorEdge, and SunSwift are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States 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. Information subject to change without notice. v3.0