An Evaluation of Sun Cluster 3.0 and SunPlex

A Foundation for Mission-Critical Applications

An IDC White Paper
Sponsored by Sun Microsystems

Analysts: Dan Kusnetzky and Jean S. Bozman

EXECUTIVE SUMMARY

Sophisticated enterprises today have multiple servers within one or more datacenters, which once acted as islands of information on their own but are now "virtually" connected through the use of clustered computing software into what appears to end users to be a single computing resource that is accessible across the network.

Service level management — the capacity to manage multiple applications across multiple tiers within a datacenter — is about trying to achieve higher service levels at lower costs. It is also concerned with reducing the risk of operation by easing the manageability of IT systems. This, in turn, increases the reliability and robustness of the total solution and addresses the customer's needs for highly available data and applications.

Clustering and high-availability software allows the data and applications that reside on these separate servers to be accessed by all end users on all of the servers within the cluster. This software also allows applications to be hosted across multiple tiers of systems within an organization's datacenter, providing high levels of robustness, availability, and manageability. More important, workloads can move from one server to another, if needed, because of resource pooling, the use of globally shared data files, and the use of storage resources provided by the virtual processing software and the underlying hardware platforms.

IDC believes that this kind of capability will become even more important in the emerging world of distributed computing and Web services. End-to-end delivery of data services — including application services that, in reality, reside on geographically dispersed servers or on multiple, physically separate servers within a corporate network — will depend upon the kind of virtualization of resources that is afforded by advanced clustering software.

5 Speen Street • Framingham, MA 01701 USA • Phone 508.872.8200 • Fax 508.935.4015 www.idc.com

In This White Paper

This white paper defines the IDC model of virtualization, clustering, and availability as a framework for understanding clustering functionality that is delivered by many vendors within the computer industry. It is organized into four major sections: IDC's view of virtual processing and clustering technology; a review of the Sun Cluster 3.0 product offerings; an IDC evaluation section, comparing Sun's clustering software to the IDC model; and a discussion of the market challenges and opportunities for Sun's clustering software.

In a fifth section, this white paper recommends that Sun's Cluster 3.0 software product and the SunPlex architectural solutions — based on Sun servers running the Solaris Operating Environment — be considered as part of an organization's IT evaluations as it builds out computing infrastructure with clustered and highly available systems. IDC concludes that Sun is a strong contender within the market space for clustered servers that deliver advanced functionality with respect to clustering technology and high-availability features.

IDC'S VIEW OF VIRTUAL PROCESSING, CLUSTERING, AND AVAILABILITY

IDC believes that clustering and virtualization are becoming increasingly important in multitier computing infrastructures, as shown in Figure 1. This is because each tier may contain multiple servers that

Figure 1: N-Tier Platform for Distributed Computing

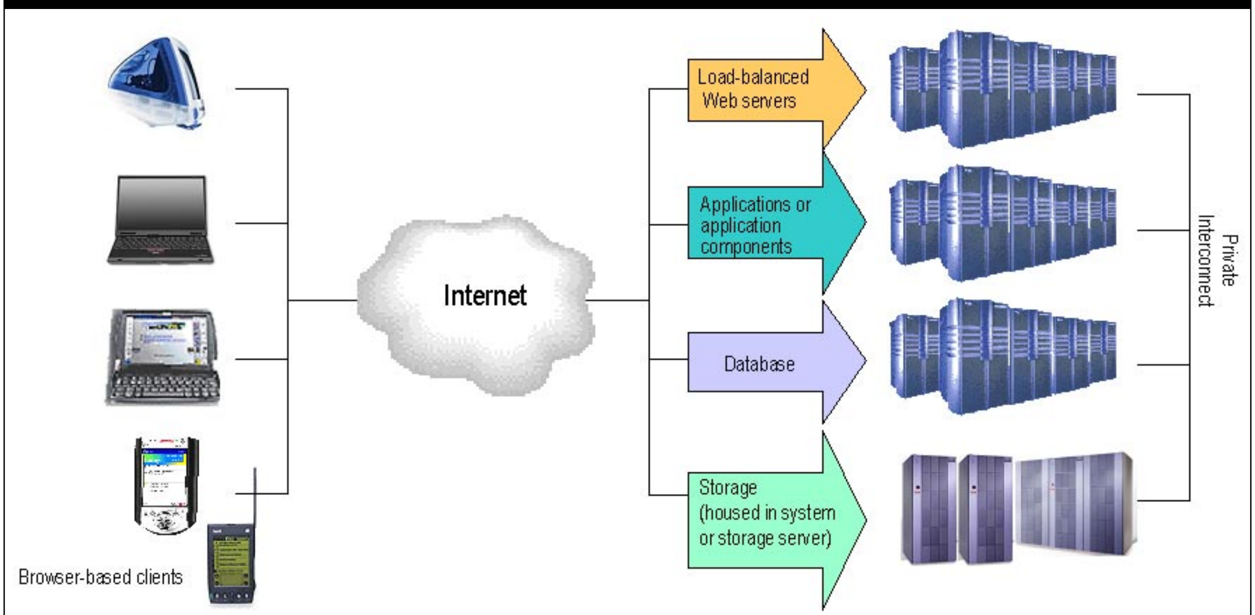

Source: IDC, 2002

can be clustered together as a virtual computing resource supporting many end users or a high volume of data requests. Alternatively, they can be clustered together to provide highly available applications and data to end users.

In the process of doing demand-side research (i.e., research about IT consumers), IDC has uncovered customer requirements for clustering and availability software and for clustering hardware. IDC has synthesized these requirements into models of virtual processing, clustering, and availability software and hardware. IDC believes that these models can serve as valuable tools for evaluating virtual processing solutions and for explaining that clustering is a combination of clustering software, server hardware, interconnects, and storage systems.

Each of the components of IDC's model of virtual processing — from software, to hardware, to interconnects, to storage — can be seen as complementary, rather than competitive, in the creation of the total clustering solution. It is clear that performance, scalability, reliability, and availability — along with the creation of a unified management domain — are the goal of organizations that are using advanced clustering and high-availability software.

Organizations are looking for these configurations to provide a reliable pool of resources from which they can draw without having to concern themselves too much with which individual server system is actually doing the work. Furthermore, these organizations want their server configurations to manage themselves by dealing automatically with day-to-day operational issues and to "heal" themselves by continuing operations, even when a failure of any individual hardware or software component occurs.

IDC's virtual processing model segments virtual processing into five system software functions providing levels of performance, scalability, reliability, availability, and manageability that exceed what is possible within a single server system alone.

These five types of system software functions include the following:

- **Performance.** Parallel processing software allows a single application to be decomposed into components or into multiple instances of an application to be run in parallel on different computers. This approach is ideal for high performance in simulation, modeling, and other applications that process large amounts of data with few interdependencies.
- **Scalability.** Load balancing software stands between the people who are using these applications and the applications themselves and distributes the workload among the systems running these applications, using either a simple round-robin approach to distribute the workload, a request/supply model in which systems notify the load-balancing software that processing capacity is available, or, in more sophisticated software, agents to sense which systems have spare processing capacity.
- **Reliability and Availability** (two categories, as shown in the IDC Model [see Table 2 in the Appendix]). Two different types of

distributed application architectures are used to ensure high availability. Storage replication software is used for data availability and application availability software, which is usually built on top of a data availability solution. Application availability software can actually restart applications. If the application has been written to use the cluster application programming interface (API), it is possible to prevent the loss of an in-process application; otherwise, it is assumed that the database will back out incomplete transactions.

- **Manageability.** Single-System Image Clustering software creates a complete virtual environment in which users see the system as a very large, very reliable, single computing resource, even though that unified virtual resource is actually constructed from a number of separate computing elements which may or may not reside in a single datacenter. A unified domain has a number of benefits: it lowers costs of administration, lowers development costs, and provides a pool of resources.

IDC's model of these five system software functions is described in detail in the Appendix.

REVIEW OF SUN'S OFFERINGS

Sun SunPlex Clustered-Server Configurations

SunPlex systems are clusters of servers that are designed to manage application services for tightly coupled computing environments. They are built around Sun Cluster 3.0 software, the Solaris Operating Environment, and Sun server, storage, and network connectivity products and services. A more in-depth definition of this configuration can be found in the section Cluster Framework Integration with the Solaris Operating Environment (see page 12).

Sun Cluster 3.0 software allows up to eight Solaris servers to be harnessed together to create what appears to the end user to be a single system that includes up to 848 processors, if eight Sun Fire 15K servers (each up to 106 processors) are used. Alternatively, customers can cluster a wide range of Sun servers — including entry servers (which, by IDC definition, are those servers priced up to $100,000), such as the Netra T1, the Sun Fire V480, or the Sun Fire V880; midrange servers (defined by IDC as those priced $100,000 to $1 million), such as the Sun Fire 4800 or 6800; and high-end servers (defined by IDC as those priced at $1 million or more), such as the Sun Ultra Enterprise 10000 or the Sun Fire 15K. Sun's classification of these three types of servers is volume servers, midrange servers, and datacenter servers (see sidebar Hardware Platforms Supported on page 5 for more examples of the Sun servers that can be included in clustered configurations).

Sun Cluster 3.0 Software

With Sun Cluster 3.0 software, application services can be categorized in one of two ways: either as a scalable service or as a failover service. The most basic difference between the two modes of operation,

DEFINITIONS

Here are some brief descriptions of the key features of Sun Cluster 3.0 software, the foundation of the SunPlex environment:

Global File Service. Abstracts the data location from the services being supplied by the cluster, delivering continuous availability of core data services, simplifying system administration, and lowering the cost of cluster deployment

Global Networking. Enables IP service to reside anywhere within a cluster environment so that service recovery is seamless to the application, which results in improved availability of network services

Global Devices. Automatically detects all storage devices upon system boot, improving ease of use, providing seamless storage scalability, and enhancing availability, regardless of the physical server node to which these devices are attached

Automated application service management. Provides virtually continuous service levels to end users

Scalable Services. Allows multiple instances of the application to run within the cluster; within scalable services, only service-level recovery is required for data processing to maintain consistent service levels

HARDWARE PLATFORMS SUPPORTED

Sun servers supported with Sun Cluster 3.0:

Entry Servers (Sun's volume servers category)

- Netra™ T1 200
- Netra 112x
- Netra 140x
- Netra 20 (Netra T4)
- Sun Enterprise™ 220R
- Sun Enterprise 250
- Sun Enterprise 420R
- Sun Enterprise 450
- Sun Fire™280R
- Sun Fire V480
- Sun Fire V880

Midrange Servers (Sun's midrange servers category)

- Sun Enterprise 3x00
- Sun Enterprise 4x00
- Sun Enterprise 5x00
- Sun Enterprise 6x00
- Sun Fire 3800
- Sun Fire 4800/4810
- Sun Fire 6800
- Sun Fire 12K

High-End Servers (Sun's datacenter servers category)

- Sun Ultra Enterprise 1000
- Sun Fire 15K

Source: Sun, July 2002

according to Sun, is the number of application instances that are executing simultaneously within the SunPlex system and also the nature of the application in the system (i.e., whether the application is "stateful" versus "stateless"). IDC notes that scalable services and failover services may both be present within the same SunPlex system.

A scalable service consists of multiple instances of the same application running simultaneously within the cluster, which provides extremely high levels of availability for a given application service. A failover service consists of only one "stateful" application instance running at any given time. If that given application goes offline for any reason, its "state" must be recovered on another Solaris domain by means of Sun Cluster's failover software capabilities. The failover process recovers the state of that application — before it went offline — and reestablishes the links between the application and the database to which the application's transactions were posted.

The Scalable Services Scenario

Sun uses the term "scalable services" to describe the support of multiple instances of a data service or application service on multiple servers — or within a single, scalable server. The workload can then be shared by a number of processing resources. This means that a larger workload can be supported by the SunPlex clustered configuration than would be possible within any single server system. This complies quite well with the scalability portion of the IDC model of virtual processing (see Table 2 in the Appendix). Readers should refer to the Evaluation section of this white paper for further analysis of the functionality of Sun Cluster's capabilities with regard to scalability.

Configuring a service as a scalable service has several benefits, such as increasing the overall scalability and availability of given applications and reducing user wait times associated with traditional clustering techniques. Since multiple instances of the application are running simultaneously, the data service does not need to be restarted if one of the instances fails.

If a failure causes one of the data services, such as an application service, to become unavailable and that data service or application service is still available somewhere else in the clustered configuration, the Global Networking capabilities of Sun Cluster 3.0 allow incoming end-user requests to be transparently rerouted to other surviving instances so that the user session is not lost. This means little or no perceived downtime. Furthermore, this process is typically faster than the traditional failover/restart approach to move data services from one server to another.

Moreover, Sun Cluster 3.0 software can be used for data-processing scalability across the multiple Solaris domains of a single, partitioned Sun server. It also can be used to link Solaris domains that are running on multiple, physically separate servers in which multiple instances of Solaris, each running within a domain, are clustered together using Sun Cluster 3.0 software and the optional Sun Management Center (Sun MC) cluster-management software.

SUN SYSTEMS-LEVEL SOFTWARE

Sun Cluster 3.0 fits into a broader world of Sun system-level software. It integrates with each of the following Sun software products:

- Solaris Operating Environment (either Solaris 8 or Solaris 9, the latest releases of the Solaris operating system, a variant of System V Release 4 [SVR4] Unix)
- Sun StorEdge Instant Image (II), storage-management software that provides data mirroring on Sun's StorEdge storage systems
- Sun StorEdge Network Data Replicator (SNDR), a data-mirroring capability that allows IT sites to ensure that mission-critical data is archived in an offsite data store, over a wide area network link, for purposes of business continuance and disaster recovery
- Sun Management Center (Sun MC), software provides customers with a control point and a unified console view of their clustered server systems; it allows system administrators to issue a command once and to have it applied to all of the servers within a single clustered configuration of multiple, individual server nodes
- Solaris Resource Manager (SRM 1.2), a software module that monitors the usage level of computing resources; works with Solaris' Automatic Dynamic Reconfiguration (ADR), which resizes partitions within a single server to complete the total solution by reassigning tasks to specific sets of available computing resources (CPU, memory, and I/O)

The High-Availability Failover Services Scenario

High-availability failover capabilities, which are shown in the IDC model (see Table 2 in the Appendix), restart applications or databases on alternative servers within a cluster, in the event that the server on which they are running goes offline or fails. Typically, restarts related to failover in clustered systems range from a minimum of 30 seconds (in the case of a server that already has the application installed) to a maximum of many minutes (in the case of a database restart). Both the failed server and the alternative server can access shared storage, which allows restart of the data service on the alternative server. In any case, the failover process itself typically entails some level of delay or interruption before end users regain access to the data service that failed.

Sun Cluster 3.0 software relies on its close integration with the Solaris Operating Environment to rapidly detect failure and to make the resultant failover sequence more reliable than might be possible through the use of layered clustering software running on top of other operating systems. The Evaluation section of this white paper provides further analysis of the functionality of Sun Cluster 3.0 with respect to failover functionality (see page 13).

Sun Cluster 3.0's failover services consist of only one instance of an application running anywhere within the SunPlex system at any given time. The failover service is most commonly used for applications in which the replication of state information is required in order for data processing to continue. When a failure of hardware or software causes the application to go offline, the application must be restarted at a new location. An end user accessing that application may experience a slight delay during the restart process. The application may be restarted elsewhere within the SunPlex clustered configuration: either in a separate partition within a scalable server or in another Sun server that is part of the SunPlex cluster.

Support for Clusterwide Resource Management

Sun Cluster 3.0-based clustered configurations offer end users a number of clusterwide resources and can be seen by system administrators from a single administrative view. Clusterwide services include Global Devices, Global File Service, and Global Network Services (which supports global IP), which are always available to the services that depend on them, even in case of failures. These services are virtualized across the SunPlex clustered-server configuration itself, and they all allow end users to access applications and data that are running on alternative compute resources within the complex in the event that any single computing resource goes offline.

In many cases, hardware or software failures are completely transparent to the applications. This is ensured by the design of Sun Cluster 3.0 software to provide direct interprocess communications (IPC) between up to eight instances of the Solaris Operating Environment, including both Solaris 8 and Solaris 9 — the new releases of the Solaris operating system, which are shipping now.

SUN CLUSTER 3.0 — A PLATFORM FOR SERVICE LEVEL MANAGEMENT

Sun Cluster 3.0 software and the SunPlex environment include the following key features:

- Global File Service that can be accessed equally by all attached servers
- Global IP services allowing network-focused applications to continue to function as long as network connectivity exists within a cluster, even when the communications subsystem of the application's host machine fails
- Systemwide service level management for system administrators to set up clusters that support quick switchover between multiple application services and data services running across the cluster; this switchover capability supports faster recovery from hardware or software faults (in any single server) than traditional failover recovery in high-availability clusters
- SunPlex Manager tool along with integration with Sun Management Center (Sun MC) for improved ease of installation, monitoring, and management of data services
- Automation of script-writing procedures for cluster enablement using the SunPlex Agent Builder tool

Flexibility of operations is supported by high-availability failover from one server to another within a cluster, by provision of alternative data services running on alternative servers, and by the notion of "scalable services" with respect to data services that are running in multiple Solaris domains throughout the cluster. Thus, clustering technology will be a key enabler of near-continuous data availability in the emerging world of Web services and provisioning of data services onto multiple computing resources. Sun Cluster 3.0 software, due to its tight integration with Solaris, appears to be able to better address this need for fast, reliable failover than might be possible with a layered software approach.

Service Level Management

Service level management is a key design point for Sun Cluster 3.0 software. Sun plans to leverage this capability in its support for virtualized computing capabilities throughout a computer network. Sun has already described its vision for the next generation of network computing, putting a new twist on the Sun motto "the network is the computer" to "the network will become the computer." The next generation of Sun systems software, including Sun Cluster software, will be designed to provide this level of computing virtualization.

Over the next several years, system administrators will start to manage virtual resources much as they currently manage today's physical computing resources. Sun's virtualization software, found in its clustering software and its storage-management software, already supports logical management of data across a single cluster. This single system provides a highly available, highly scalable platform for applications.

Sun Cluster 3.0 software leverages the capabilities that Sun has built into its midrange "midframe" servers, such as the Sun Fire 6800, which supports up to four partitions; the Sun Fire 12K, a midrange server priced from $500,000, which supports up to nine partitions; and the Ultra Enterprise 10K and Sun Fire 15K high-end servers, in which up to 18 partitions (or domains) can be created, each running a separate instance of Solaris. When multiple Solaris instances are running on a single machine, it is possible to use Sun Cluster 3.0 technology to move applications from one domain to another. Using Solaris' built-in Automatic Dynamic Reconfiguration (ADR) capability, it is also possible to dynamically resize the domains within a Sun cluster, as needed, according to the compute requirements of the application workloads that are running.

Sun server customers can now set up multiple domains on single servers, or they can set up those domains across multiple servers.

Some IT managers find it advantageous to set up these domains for purposes of server consolidation and high availability of mission-critical data. When that is done, Sun Cluster 3.0 and its Sun Management Center (Sun MC) management capabilities can push workloads to the appropriate compute resources within a single Sun server.

Alternatively, multiple Solaris domains can be set up across multiple, physically separate servers within a cluster. This approach reduces the impact of any one server going offline because alternative Solaris domains running on other servers are immediately available for switchover within the cluster.

IDC believes that access to all attached server nodes within a Sun Cluster 3.0 configuration can be compared to the type of software virtualization that was previously seen in traditional, high-end, host computer systems, which provide the highest levels of reliability and data availability. For example, IBM-compatible mainframes running in an IBM SysPlex made up of many servers, each of them running an IBM operating system (including the IBM z/OS, the OS/390 operating system, the IBM VM/XA operating system and its follow-on z/VM virtualization technology, or the IBM TPF operating system) could provide similar levels of uptime to that provided by Sun's solution. Alternatively, similar configurations could be formed through the use of multiple Compaq Himalaya high-end server systems, which could achieve large single system images by clustering many individual server nodes together (see Table 1 in the Appendix). Sun is bringing this kind of datacenter host functionality to the world of distributed Unix servers that are rapidly taking on many of the mission-critical workloads once reserved for mainframes.

In the future, however, computing resources will be managed at an even higher level of abstraction than is possible today — involving multiple clusters and geographically distributed servers. Thus, Sun Cluster 3.0 was built as part of a technology road map for wider virtualization that will play out over the next five years, IDC believes. In that respect, Sun Cluster 3.0 can be seen as one of the building blocks of a wider Sun technology foundation for the next wave of distributed computing, which aims to deploy data services across a wide array of networked servers and to manage these data services at a high level of abstraction — rather than to manage them on a server-by-server basis, as is the practice in today's datacenters.

Data Availability

Upon system boot, Sun Cluster 3.0 detects all storage devices — other than its direct-attached storage — through its Global Devices feature. All of these storage devices are available to any application on any system in the cluster. This allows cluster storage to appear to be a single resource even though the storage devices may be physically attached to one of the cluster nodes, attached to a storage server and being accessed via a storage area network, or network attached storage being accessed by one member of the cluster.

Global Devices make storage easier to use, simplify cluster administration, and make it possible for system configurations to include redundant storage hardware. This approach means that failures of storage volumes, storage adapters, or storage cables can be hidden from applications. If a device fails or becomes inaccessible for some reason, the application can continue processing using a replica of that volume.

Global file service virtualizes storage so that data no longer needs to be physically attached to the server that is hosting the application or database.

Global network service virtualizes access to network interfaces. This means that a network application service can actually be hosted anywhere within the cluster. Multiple IP addresses can be used for a single service, or a single IP address can be used to access multiple services. Sun's Global Network Services creates a many-to-many relationship between network addresses and network application services.

This also means that applications can continue to run even though individual network interfaces have been stopped or restarted. Network application services can be stopped and restarted on another system, without making that service unavailable, through the use of Sun Cluster 3.0's built-in workload balancing capability.

So, the failure of a network adapter or network application service no longer has to mean that users will experience an interruption. Furthermore, the built-in, load-balancing mechanisms would also tend to better utilize overall cluster resources.

Application Availability

Applications or application components can be hosted by several systems within a cluster. This capability can be used in several ways: reliability, manageability, performance, and scalability. Multiple instances of a given application can be run simultaneously and each instance fed different portions of a single data set to provide higher levels of performance. Another use of this capability is extending system scalability to allow more user workloads or larger data sets to be processed.

Sun Cluster 3.0's failover services allow multiple mission-critical or business-critical applications and services to share the same clustered hardware configuration. This failover capability in and of itself is comparable to that of other Unix systems vendors.

In addition, Sun provides the scalable services function, which leverages workload-balancing software to direct the incoming end-user requests to other instances of those applications or network application services that are running elsewhere on the cluster. For example, this capability is already supported for SAP applications running on multiple Solaris instances within a SunPlex system. By building up these capabilities, Sun is taking strides in its progression toward building a "services-driven network" in which "data services" are provided via provisioning of those data services onto several servers within a given corporate network, intranet, or extranet.

To enable these scalable services, Sun has provided Sun Cluster 3.0 with the capability to quickly detect an error and to invoke the workload-balancing software to move end-user requests to the alternative data service. In that scenario, users would see only a short pause or interruption in service when such a failure occurs. This improves upon the situation in many clustered server configurations throughout the industry, in which failover times for specific applications or databases may take several minutes to be reestablished on other servers

SUPPORT FOR VIRTUALIZATION OF THE SERVER RESOURCES

The Sun Cluster 3.0 software, combined with Sun's SPARC/Solaris server systems, storage devices, and systems management software, offers organizations an application service level management platform that delivers overall levels of performance, scalability, application availability, and data availability — as well as the ability to create a unified management and computing domains from a number of physically separate systems. Sun Cluster 3.0 is the key element of Sun's SunPlex computing environment, which includes the hardware, software, and storage components of the overall solution.

The Sun Cluster 3.0 capabilities allow IT organizations to consider their clustered Sun systems to be a unified pool of virtual resources that allows them to establish a foundation for managing the delivery of mission-critical and business-critical data services — all of which are key requirements for the next wave of enterprise computing. IDC believes that these capabilities will take on even more importance as the next generation of Web services starts to be delivered over the next two to three years.

within the cluster. More important, if a failure occurs somewhere within the cluster, applications or network application services running on other cluster nodes continue to run. Users of those applications or services would be unaware of a failure.

This is comparable to the situation in which multiple applications are running on a series of servers in a "server farm" in a multitier or n-tier environment (refer back to Figure 1). In that case, alternate data services are available as soon as the end user can connect to a new session with the alternate application server. In the case of databases, the situation is a bit different, in IDC's view, because more than one database instance would have to be synchronized with a mirror copy of that database in order to retain comparable data access in case one of the two databases went offline.

However, IDC notes that Sun Cluster 3.0's load-balancing facilities allow a large workload to be shared among several Sun systems, improving overall scalability and performance.

Application Management

Sun Cluster 3.0's Resource Group Manager (RGM) takes the concept of service level management one level higher by allowing groups of applications, application components, and system services to be seen as a single resource. Applications and application components can be registered with the RGM component of the cluster framework, allowing administrators to group functions together. This allows the framework to respond intelligently to failures. The cluster can reconfigure itself to keep registered applications running without affecting applications or system services that do not rely on the failed functions. The entire group of applications can be quickly rehosted in the case of a failure, without requiring operator intervention. This type of rapid deployment capability helps to ensure near-continuous availability of mission-critical applications and data.

Virtualization Offered by Sun Cluster 3.0 Software

The virtualization provided by the combination of Sun's Global Devices, Global File Service, Global Networking, and Scalable Services features provides a quick restart capability in the event that the system initiates the failover of a single application, since the file system and the cluster-attached storage devices would not have to be restarted. This dramatically reduces the service-interruption time for end users.

Thus, mission-critical and business-critical applications will remain available to end users on a continuous or near-continuous basis, depending on whether a failure occurs in any given server. IDC notes that the more domains that have been set up within a Sun cluster (either with multiple servers or within a single, scalable server's domains), the less the likelihood that end users would experience appreciable disruption to data access.

The Sun Cluster 3.0 framework presents administrators with a unified management domain. This means that all of the cluster resources can be managed and administered as if they were on a

SOFTWARE AGENTS

List of agents supported on Sun Cluster 3.0:

- Apache Proxy Server
- Apache WebServer (scalable and HA)
- DB2
- iPlanetTM Web Server (scalable and HA)
- iPlanet Messaging Server
- iPlanet Directory Server
- iPlanet Calendar Server
- Oracle Enterprise Server
- Oracle Parallel Server (OPS)
- Oracle 9i RAC
- Informix
- SAP (Scalable)
- SBU
- Sybase
- Netbackup
- Broadvision
- Sun DNS
- Sun NFS

Source: Sun, July 2002

single system. These resources can be viewed through the Sun Management Center 3 console and via browser-based Web connections, from remote sites such as a system administrator's home. This capability allows a single system administrator to become much more effective and productive, as any action or command needs to be executed only once to be applied throughout the cluster. Clustering software packages in the industry that do not yet support a global file service (along with smooth interprocess communications between operating systems across the cluster) require that system administrators apply each clusterwide command repeatedly to each server within a given cluster. This is a time-intensive approach that reduces system administrator productivity, compared with deploying commands simultaneously throughout the cluster.

Application Building with Sun Software Tools

The Sun Cluster 3.0 API enables off-the-shelf applications to be integrated into the cluster framework. This broadens the reach of Sun Cluster use to a wide variety of packaged applications and packaged databases. Furthermore, a software development kit (SDK) and agent-builder wizard tools provided in Sun's SunPlex Agent Builder toolkit are available to speed the development effort required to create cluster-aware applications. It also allows an agent or cluster to become application-aware. Thus, Sun says that application services are enabled by the combination of the application itself and the Sun Cluster agent. This is particularly useful for custom applications written for a specific industry or enterprise, since it allows those applications to be made cluster-aware more quickly and efficiently than if IT staff programmers had to develop one-off clustering scripts for each custom application.

Since its introduction in December 2000, the Sun Cluster Agent Builder tool has allowed customers to create both failover and scalable agents rapidly. This fast-scripting capability complements the provision of predefined cluster-aware software agents that are provided by Sun and its independent software vendor (ISV) partners. Among these predefined agents are those for the Oracle, Informix, and IBM database products, along with those for SAP, Sybase, Oracle 9i RAC, and the open source-based Apache Web Server.

SunPlex Agent Builder works by generating "agent" code in the Korn Shell's ksh or C languages, with automation of the scripting process that is triggered by means of a graphical user interface (GUI)-based start command. In December 2001, Sun enhanced its SunPlex Agent Builder so that it now supports the Generic Agent functionality that generates a precompiled agent binary without any further programming or modification to existing code.

To ensure high levels of security, programmers can use the Solaris Security Toolkit, which is also known as the JumpStart Architecture and Security Scripts (JASS) toolkit. JASS provides a flexible and extensible mechanism to minimize, harden, and secure systems that are based on the Solaris Operating Environment system. The primary goal behind the development of this toolkit is to simplify and automate the process of securing Solaris systems.

Cluster Framework Integration with the Solaris Operating Environment

Sun Cluster 3.0 technology is integrated into the Solaris Operating Environment. It virtualizes the Solaris Operating Environment, allowing applications be both highly available and highly scalable by allowing multiple application "instances" to run across the Sun clustered-servers configuration. To achieve this, Sun leveraged five years of research and development work that resulted in the ability of Sun Cluster 3.0 to optimize interprocess communication (IPC) between multiple instances of the Solaris operating environment. This work, done as part of Sun's "Full Moon" development project in Sun Labs, resulted in differentiated Sun technology that confers high levels of virtualization for servers that are participating in a Sun Cluster 3.0-clustered configuration.

The Sun Cluster 3.0 framework enables resources, such as devices, file systems, and networks, to be shared in the cluster. It provides a single, integrated hardware and software platform for applications that share the available cluster resources. More resources can be added to the platform by adding more resources to the cluster, either in the form of additional cluster nodes or by adding resources to existing nodes.

Sun SunPlex Clustered-Server Configurations

SunPlex systems are clusters of servers that are designed to manage application services for tightly coupled computing environments. They include the following hardware and software components:

- Sun Cluster 3.0 software
- Two or more instances of the Solaris 8 or Solaris 9 server operating systems
- Sun Cluster 3.0 supports the whole range of the Sun server line, starting from the one-way Netra T1 to the 106-way Sun Fire 15K. In addition, Sun Cluster configurations can include the Enterprise 250, 450, 220R, and 420R entry servers and the Sun Enterprise 3000, 3500, 4500, 6000, and 6500 midrange servers as long as these UltraSPARC II-based servers run Solaris 8 or Solaris 9. All of the Sun Fire servers, which are based on UltraSPARC III microprocessors and the Solaris 8 or Solaris 9 operating systems, can also be deployed as server nodes for SunPlex configurations. These include Sun's "Mid-Frame" Sun Fire 3800, 4800, and 6800 servers, the Sun Fire 12K high-end midrange server, and the Sun Fire 15K high-end server.
- Packaged software modules for applications run in a high-availability mode across the clustered servers. Alternatively, customers can use Sun scripting toolkits to create their own cluster-aware applications for the SunPlex system. The Sun scripting toolkits are designed to reduce overall development time through automation of the coding process. The Sun Cluster API and the SunPlex Agent Builder (see the section on page 11) provide options to facilitate packaged software's ability to make full

PRECONFIGURED CLUSTERS

Sun markets several preconfigured clusters as a way to reduce the time for deployment, setup, and installation.

The Sun hardware, Sun system software, and Sun clustering software are preinstalled and preassembled into an integrated "cluster in a rack" in the Sun factory.

At the same time, the use of preconfigured clusters reduces the need to have IT staffers with advanced HA skills at each deployment site, thus lowering the IT skills requirement for customer sites that wish to deploy a Sun clustered-server configuration. All preinstalled software is tested and certified by Sun Engineering.

Cluster Platforms from the SunTone™ Platform product portfolio are integrated stacks of Sun hardware and software designed to provide a robust, high-availability framework with failover protection and built-in workload-balancing capabilities.

Here is a brief list of the predefined clusters, which will be extended in the future to include other sets of servers and storage systems:

- Cluster Platform 220/1000, which includes two Enterprise 220 servers, both attached to shared disks within one or more Sun 1000 storage systems
- Cluster Platform 4500/3, which includes two Enterprise 4500 servers, both attached to one or more T3 storage systems
- Cluster Platform 280/3, which includes two Sun Fire 280R servers, attached to one or more T3 storage systems
- Cluster Platform 15K/9960, which includes two Sun Fire 15K servers, attached to one or more StorEdge 9960 storage systems

use of the clustered configuration. The use of the API confers cluster-awareness on any packaged or custom application. The Agent Builder speeds the development of code through automation of the coding process itself.

- Sun StorEdge storage arrays, or third-party storage devices provided by other vendors, include the Sun T3 storage device, the StorEdge D1000, StorEdge 3900, StorEdge 6900, and StorEdge 9900 storage systems. In addition, third-party storage devices, including those from EMC Corp., are also supported.
- Hardware interconnects between all of the clustered servers and the network connectivity (locally, in a campus, or via a wide area network) create the clustered-server configuration. Typical interconnects include Fast Ethernet (100Mbps Ethernet), Gigabit Ethernet (GbE), and SCI over the Remote Shared Memory (RSM). Sun also provides a Campus Clustering solution that connects clustered servers via a long-distance, wide area network (WAN) connection. Additionally, Sun sells the Sun StorEdge Network Data Replicator (SNDR), a mirroring facility that supports business continuance and disaster-recovery planning by enabling mirroring via geographic clusters, in which the clustered servers are connected via a WAN link. It also sells Sun StorEdge 9900 TrueCopy Software with its StorEdge 9900 high-end storage systems.

The design point for the SunPlex systems, whether preconfigured by Sun at the factory or configured and installed at a customer site, is to create a single, virtual computing resource that improves the reliability, scalability, and manageability of the core system and ease of use for end customers.

Deployment of SunPlex clustered-server configurations is intended to reduce the customer's total cost of ownership (TCO) over time, even as it provides such near-term benefits as high availability for mission-critical and business-critical applications, support for server consolidation, reduced management requirements for system administrators, and reduced system downtime for end users.

The provision of preconfigured configurations reduces setup and installation time and allows customers to adopt optimized configurations rather than spends time developing custom scripts and requiring system administrators to spend time learning how to install specific combinations of servers and storage systems. IDC research has shown that many customers delay installation of enterprise clustered systems because of concerns about having sufficient in-house IT experience to deploy such clusters and about spending time developing scripts to make custom and package applications "cluster-aware."

EVALUATION: HOW DOES SUN'S CLUSTERING SOLUTION COMPARE TO IDC'S MODEL?

As Table 1 shows, the Sun SunPlex Cluster 3 solution clearly fits into the clustering software segment of IDC's model of clustering and availability software. It goes beyond that segment in several ways by providing higher levels of virtualization, allowing the configuration to be seen as a large pool of resources rather than a number of individual systems.

Table 1: Evaluation of Sun Cluster 3.0

Feature/Function	Sun Cluster 3.0
Data availability	The virtualization provided by the combination of Sun's Global Devices and Global File Service means that the failure of any single component of the storage subsystem would not cause data to become unavailable. Users would experience little or no service interruption time.
Application availability	The virtualization provided by the combination of Global File Service and Scalable Services means that applications will appear not to fail. Sun Cluster 3.0's Resource Group Manager takes this to the next level of virtualization by acknowledging that applications often consist of many components and system services that may be hosted on different systems within the cluster (see sidebar Software Agents on page 11 for supported applications).
Parallel performance	Sun's scalability services mean that organizations can run multiple instances of applications or application components. The result would be higher levels of performance than those possible with a single system. Sun Cluster 3.0's SDK and Agent Builder tool allow applications to be built to enhance parallel performance.
Scalability	Organizations can support more users or larger databases by running multiple instances of applications or application components. The result would be higher levels of scalability than possible with a single system.
Unified domain	Sun Cluster 3.0 presents a single system image to everyone, reducing the perceived complexity of the environment and the related costs due to complexity. Sun Cluster 3.0 takes the next step by presenting an environment in which multicomponent applications can be seen as, managed as, and restarted as a single application. Sun's Global File Service and Global Networking Services are key features creating a unified domain.

Source: IDC, 2002

The addition of Sun's SunPlex (a clustered configuration of Sun servers — which can include high-end servers, midrange servers, or entry-server systems), Sun's storage software, and hardware allows Sun to fulfill all of the requirements of all of the different types of clustering and availability architectures.

Addressing the Scalability and Availability of Data Services

IDC notes that Sun Cluster 3.0's scalable services are useful for increasing both overall system scalability and the availability of business-critical and mission-critical applications. Operation in the scalability mode does not meet the IDC model's requirements for a failover clustering solution that supports the failover of "stateful" applications. This is because any transactions not updated before a stateful application goes offline would have to be recovered from a system log and then resubmitted to an alternate data service for completion. However, in the event that stateless applications (e.g., support of static Web sites, firewalls, and security/authentication software) were referred to an alternative Solaris domain for processing, then the scalability services would meet the criteria of providing very high levels of availability for those data services — and for the end users accessing those data services.

With regard to Sun Cluster 3.0's failover services, these meet the criteria for failover high-availability clustering for stateful applications, according to the IDC model. IDC believes that Sun Cluster 3.0 provides advanced functionality that competes strongly with other failover clustering technologies in the marketplace. Sun Cluster 3.0's Global File Service, with its support for clusterwide virtualization, is in the top tier of high-availability failover software packages shipping in the market today (see Table 1). IDC believes that when both the failover services and scalability services are used, Sun Cluster 3.0 provides a very comprehensive approach to clustering and high availability in both the scalability and failover dimensions.

CHALLENGES AND OPPORTUNITIES

Challenges for Sun IT infrastructure within many organizations often consists of applications, application services, and network services that run on a variety of servers made by different vendors. IDC research has shown that most large organizations have deployed multiple types of servers, running multiple types of operating systems, to run multiple workloads in heterogeneous, mixed-vendor computing environments. Furthermore, large corporate networks have multiple domains — and each of these domains supports collections of computing workloads. Sun must find a way to confer the benefits of its Sun Cluster 3.0 technology on this mixed-vendor, heterogeneous world of computing, whether as a foundation for datacenter hosting, as an anchor for corporate data networks, or as a data-management hub within a diverse array of heterogeneous systems running on top of cross-platform Java/XML middleware.

Sun has a vision for building an enterprise management framework, called N1, which will manage application services that cross multiple computing domains. Sun's clustering software should fit within this N1 framework, both as the virtualization engine for clustered-server configurations and as a means to link multiple partitions within a single, scalable Sun Fire 15K server. However, IDC believes that it will likely take several years for Sun to fully productize its N1 vision with multiple virtualization engines that link resources running on different types of servers (including both Sun servers and nonSun servers).

In the competitive landscape for Sun's clustering and virtualization software, each of the major systems vendors — IBM, HP, and Sun — has identified the opportunity of helping IT to manage workloads and application services across heterogeneous, mixed-vendor networks. All three have clustering and virtualization software to bring to bear on this new set of computing challenges. All have embarked on this journey to cross-platform management of computing "fabrics" — and large corporations will likely install software products from more than one of these vendors in some or all of their major sites. Thus, Sun and its competitors will need to support interoperability and open networking standards to acknowledge that multiple solutions will likely coexist in these next-generation computing "fabrics."

CONCLUSIONS

Sun Cluster 3.0 is a highly functional clustering software product, as can be seen in this white paper. Customers should compare its functionality with that of the wide spectrum of clustering products on the market today. Sun has worked hard to advance the state of the art for clustering technology, especially in the area of providing high levels of application availability and data availability for mission-critical line-of-business (LOB) operations. The ability to manage the application services that are running across the cluster in a virtualized environment should go a long way toward moving computing infrastructures, including Sun systems, toward a new era in which application services are provided and managed across a network of interconnected servers and storage systems.

It is clear that Sun's clustered-server architecture can fulfill the requirements of an organization's needs for virtual processing. Access to all attached server nodes within a Sun Cluster 3.0 configuration can be compared to the type of software virtualization that was previously seen only in traditional high-end host computer systems, which, IDC believes, provide the very highest levels of reliability and data availability.

CONSIDERATION FOR IT EVALUATION OF SUN CLUSTER 3.0

As can be seen in this paper, the Sun Cluster 3.0 software and the clustered-server configurations that can be created with it satisfy the key hardware and software requirements for clustering and high-availability technology. Applications can be decomposed to run in parallel to provide enhanced performance. Multiple instances of a single application can be run to provide enhanced scalability and/or enhanced reliability. A benefit to Sun's approach is that applications need not be rearchitected or rewritten.

On this basis, IDC believes that Sun's cluster-related products should be considered by organizations having the following requirements:

- The need for a single application to execute faster than is possible with the fastest available single system
- The need to support more users than is possible with any currently available single system
- The need for applications and data to offer higher levels of availability and reliability than can be accomplished using a single system
- The need to provide a unified management domain for multiple workloads that are running in multiple partitions — either within a single server or in multiple servers within a cluster

As we have seen in this white paper, Sun Cluster 3.0 is solidly in the top tier of clustering-software products in the marketplace. Its support for advanced functionality, its ability to share files across the cluster, and its support for operational flexibility via the scalable services feature will both support customers' efforts to improve their TCO and make their daily computing operations more efficient.

APPENDIX A

Virtual Processing Model

IDC's model of virtual processing segments virtual processing into five system software functions — providing performance, scalability, reliability, and availability beyond that possible by using a single system while still presenting a unified management domain. These five system-software functions are described in Table 2.

Enhanced Performance

Parallel processing software allows a single application to be decomposed into components or into multiple instances of an application to be run in parallel on different computers. The parallel processing software feeds data to these application components or applications, then gathers up the results and consolidates the results into a single data set. The net result is that data is processed faster than would be possible using a single machine.

Raw performance, not application availability, is the goal of this distributed processing approach. No attempt is made to ensure that applications never fail. If an application component or application instance fails for any reason, the segment of the input data set, which was being processed, is simply sent to another instance of the application component or application.

This approach is often used for simulation, modeling, or processing large amounts of data having few interdependencies. Many commercial applications process data having many interdependencies. This means that this approach to distributed processing is not applicable.

Enhanced Scalability

This distributed processing architecture deploys multiple instances of applications so that they can be processed on multiple servers. Load-balancing software stands between the people who are using these applications and the applications themselves. This software is sometimes referred to as load balancing, quality of service (QoS) software, or several other vendor-specific names.

Regardless of the name used to describe the function, as requests come in, the workload is distributed among the systems running these applications. Some load-balancing software products use a simple round-robin approach to distributing the workload. Others use a request/supply model in which systems notify the load-balancing software that processing capacity is available. More sophisticated products use agents to sense which systems have spare processing capacity, and the workload is sent to these systems.

Table 2: Model of Clustering and Availability Software

Feature/Function	Parallel Processing Software	Load-Balancing Software	Availability Software	Clustering Software
Philosophy or concept	Raw performance	Raw scalability	Increased availability of applications and/or data	Single system image
Data availability	Not a primary design goal. Data remains available through over-provisioning. Multiple copies of data are maintained on separate systems.	Not a primary design goal. Data remains available through over-provisioning. Multiple copies of data are maintained on separate systems.	Data remains available through over-provisioning. Multiple copies of data are maintained on separate systems.	Data remains available through clusterwide file system, which provides mirroring or replication capabilities.
Application availability	Not a primary design goal. Applications remain available through over-provisioning. Messages are passed to surviving applications or application components.	Not a primary design goal. Applications remain available through over-provisioning. Messages are passed to surviving applications or application components.	Applications are rewritten to use cluster APIs provided by the high-availability (HA) manager. The HA manager can then track the progress of the application and restart it on another cluster node if necessary. Non–cluster-aware applications can be restarted using cluster scripts.	Applications can be restarted using scripts. Only applications requiring access to special cluster features must be rewritten to use the cluster APIs provided by the cluster manager.
Parallel performance	Single-application performance improvements created by decomposing the application into components and running the components separately on different machines. Data sharing between components is difficult to impossible.	Not a primary design goal. Single-application performance improvements created by decomposing the application into components and running the components separately on different machines. Data sharing between components is difficult to impossible.	Not a primary design goal. Single-application performance improvements are created by decomposing the application into components and running the components separately on different machines. Data sharing is possible by rewriting the application to use special HA manager APIs.	Single-application performance improvements are created by decomposing the application into components and running the components separately on different machines. Data sharing is possible using a clusterwide file system without requiring that the application be rewritten.
Scalability	Not a primary design goal. Application scalability is through over-provisioning. Data sharing is difficult to impossible.	Application scalability is through over-provisioning. Data sharing is difficult to impossible.	Not a primary design goal. Application scalability is through over-provisioning. Data sharing is possible by rewriting the application to use special HA manager APIs.	Application scalability is through over-provisioning. Data sharing is possible using clusterwide file systems without requiring that the application be rewritten.
Management	Not a primary design goal. Systems appear to be separate, distinct management and development domains. Configuration appears to be only a single resource to the end user.	Not a primary design goal. Systems appear to be separate, distinct management and development domains. Configuration appears to be only a single resource to the end user.	Not a primary design goal. Systems appear to be separate, distinct management and development domains. The configuration appears to be a single resource to the end user, but tools exist to allow all systems to be managed together.	Systems appear to be a single computing resource. Administrators, developers, and users can see the configuration as a single machine. Users can submit batch jobs and not know where they run. Users can submit print jobs and not know which system owns the printer. Developers write applications the same way as if they were going to run on a single system. Users have no idea which system is supporting their workload.
Example products	Sun Gridware, TurboLinux EnFuzion, Platform Computing LSF	Microsoft Application Center 2000, Red Hat High Availability Server, TurboLinux Cluster Server	Legato Cluster Server, HP MC/ServiceGuard, IBM HACMP, Microsoft MCSC, Mission Critical Linux Convolo, Veritas Cluster Server	Sun Cluster 3.0, Compaq TruCluster, Caldera Non-stop Clusters for OpenUnix, IBM Sysplex for z/OS

Source: IDC, 2002

As with parallel processing software, application availability is a side effect rather than the primary goal. The primary goal of this software is scalability. The assumption is that the database will back out intermediate results. The users are expected to re-enter any data that was not retained following a server outage. More sophisticated applications store the data users have entered as cookies on the client system or in temporary files on the server. This data is retrieved when the user restarts the transaction.

If an application environment has failed, the next application request is simply routed to another instance of the application.

Enhanced Reliability and Availability of Applications

Application and data availability software is often selected when the workload requires that only the data and application files be virtualized. The goal is making both the data and the applications highly available and reliable.

Data Availability Software

Two different types of distributed application architectures are used to create a highly available environment. Data availability is provided by storage replication software. This software maintains one or more replicas or mirrors of files, directories, file systems, or complete volumes. If a storage adapter, cable, storage volume, or some other component of a storage system fails, this data availability software simply sends the request to another replica.

Highly available data, not applications, is the goal of this type of software. Unless applications have been specially written to use cluster APIs, they will be unaware of the failure and unable to respond to it — making it possible for a transaction to be lost. It is assumed that the database or application software is maintaining the state of the application and will back out incomplete transactions.

Application Availability Software

This software is usually built on top of a data availability solution and actually restarts applications. If the application has been written to use the cluster API, it is possible to prevent the loss of an in-process application. Otherwise, it is assumed that the database will back out incomplete transactions.When the application restarts, the user will be asked to start over again.

Single-System Image Clustering Software

Clustering software, unlike all of the distributed processing software described earlier, creates a complete virtual environment. Nearly everyone — users, developers, and administrators — believes that they are working with a large, reliable, single computing resource, even though that unified virtual resource is really constructed from a number of separate computing elements which may or may not reside in a single datacenter.

Applications can receive some benefit from this environment (automatic restart) without being rewritten to use cluster APIs. To receive full benefit of the environment, such as being informed when something fails so that specific actions can be taken to facilitate restart without data loss, however, still requires that applications be written to use cluster APIs to obtain cluster state information. Right now, only four different software products have the ability to create a single system image.

Unified Domain Lowers Cost of Administration

By offering a unified environment, clustering software can significantly lower an organization's cost of administration (system, network, and application). One or a small number of administrators can monitor the virtual system and move applications or application components around as necessary to maximize performance. This approach can also lower an indirect cost: the cost of lost productivity when an application, application component, or network service becomes unavailable.

Unified Domain Lowers Cost of Development

Clustering software offers all of the functions offered by parallel processing software, load-balancing software, data availability software, and application availability software, and it also creates a single system view. Administrators need know only one cluster management framework. Developers see a single development environment. Piecemeal solutions constructed organically (growing outward from a core of one-off components) tend to be more expensive when compared with a carefully architected solution.

Unified Domain Provides Pool of Resources

Some clustering software goes beyond a single system image by creating a higher level of virtualization — sometimes referred to as QoS, service level management, or even class scheduling. Rather than simply pulling tasks off of a queue using a simple priority-based scheme, added intelligence allows the system resources to be allocated by a number of factors, including the following:

- **Date and time.** Resource-intensive functions can be held in the queue until "after hours" and then requeued for later processing when normal business hours arrive. Some tasks could be held for week-end or month-end processing, and so on.
- **Task importance.** Tasks deemed to be important could be done first rather than simply waiting their turn. This function allows organizations to provide response time guarantees in response to a service level agreement.
- **Requestor.** Tasks from specific departments or individuals are executed first.

Although some operating systems provide this type of virtualization on a single system, providing this virtualization within a clustered environment is a critical component if the cluster is to be seen as a pool of managed resources.

IDC Worldwide Offices

CORPORATE HEADQUARTERS

IDC
5 Speen Street
Framingham, MA 01701
United States
508.872.8200

NORTH AMERICA

IDC Canada
36 Toronto Street, Suite 950
Toronto, Ontario M5C 2C5 Canada
416.369.0033

IDC California (Irvine)
18831 Von Karmen Avenue
Suite 200
Irvine, CA 92612
949.250.1960

IDC California (Mountain View)
2131 Landings Drive
Mountain View, CA 94043
650.691.0500

IDC New Jersey
75 Broad Street, 2nd Floor
Red Bank, NJ 07701
732.842.0791

IDC New York
2 Park Avenue
Suite 1505
New York, NY 10016
212.726.0900

IDC Texas
100 Congress Avenue
Suite 2000
Austin, TX 78701
512.469.6333

IDC Virginia
8304 Professional Hill Drive
Fairfax, VA 22031
703.280.5161

CENTRAL AND EASTERN EUROPE

IDC CEMA
Central and Eastern
European Headquarters
Male Namesti 13
110 00 Praha 1
Czech Republic
420.2.2142.3140

IDC Croatia
Srednjaci 8
1000 Zagreb
Croatia
385.1.3040050

IDC Hungary
Nador utca 23
5th Floor
H-1051 Budapest. Hungary
36.1.473.2370

IDC Poland
Czapli 31A
02-781 Warszawa, Poland
48.22.7540518

IDC Russia
Suites 341-342
Orlikov Pereulok 5
Moscow, Russia 107996
7.095.975.0042

MIDDLE EAST AND AFRICA

IDC Middle East
1001 Al Ettihad Building
Port Saeed
P.O. Box 41856
Dubai, United Arab Emirates
971.4.295.2668

IDC Israel
4 Gershon Street
Tel Aviv 67017, Israel
972.3.561.1660

IDC South Africa
c/o BMI TechKnowledge
3rd Floor
356 Rivonia Boulevard
P.O. Box 4603
Rivonia 2128, South Africa
27.11.803.6412

IDC Turkey
Tevfik Erdonmez Sok. 2/1 Gul
Apt. Kat 9D
46 Esentepe 80280
Istanbul, Turkey
90.212.275.0995

EUROPE

IDC Austria
c/o Loisel, Spiel, Zach Consulting
Mayerhofgasse 6
Vienna A-1040, Austria
43.1.50.50.900

IDC Benelux (Belgium)
Boulevard Saint Michel 47
1040 Brussels, Belgium
32.2.737.76.02

IDC Denmark
Omøgade 8
Postbox 2609
2100 Copenhagen, Denmark
45.39.16.2222

IDC Finland
Jarrumiehenkatu2
FIN- 00520 Helsinki
Finland
358.9.8770.466

IDC France
Immeuble La Fayette 2
Place des Vosges Cedex 65
92051 Paris la Defense 5, France
33.1.49.04.8000

IDC Germany
Nibelungenplatz 3, 11th Floor
60318 Frankfurt, Germany
49.69.90.50.20

IDC Italy
Viale Monza, 14
20127 Milan, Italy
39.02.28457.1

IDC Netherlands
A. Fokkerweg 1
Amsterdam1059 CM, Netherlands
31.20.6692.721

IDC Portugal
c/o Ponto de Convergancia SA
Av. Antonio Serpa 36 - 9th Floor
1050-027 Lisbon, Portugal
351.21.796.5487

IDC Spain
Fortuny 18, Planta 5
28010 — Madrid
Spain
34.91.787.2150

IDC Sweden
Box 1096
Kistagangen 21
S-164 25 Kista, Sweden
46.8.751.0415

IDC U.K.
British Standards House
389 Chiswick High Road
London W4 4AE United Kingdom
44.208.987.7100

ASIA/PACIFIC

IDC Singapore
Asia/Pacific Headquarters
80 Anson Road
#38-00 IBM Towers
Singapore 079907
65.6226.0330

IDC Australia
Level 3, 157 Walker Street
North Sydney, NSW 2060
Australia
61.2.9922.5300

IDC China
Room 611, Beijing Times Square
88 West Chang'an Avenue
Beijing 100031
People's Republic of China
86.10.8391.3610

IDC Hong Kong
12/F, St. John's Building
33 Garden Road
Central, Hong Kong
852.2530.3831

IDC India Limited
Cyber House
B-35, Sector 32, Institutional
Gurgaon 122002
Haryana India
91.124.6381673

IDC Indonesia
Suite 40, 17th Floor
Jakarta Stock Exchange
Tower 2, Jl. Jend. Sudirman Kav. 52-53
Jakarta 12190
6.221.515.7676

IDC Market Research (M) Sdn Bhd
Jakarta Stock Exchange Tower II
17th Floor
Jl. Jend. Sudirman Kav. 52-53
Jakarta 12190
62.21.515.7676

IDC Japan
The Itoyama Tower 10F
3-7-18 Mita, Minato-ku
Tokyo 108-0073, Japan
81.3.5440.3400

IDC Korea Ltd.
Suite 704, Korea Trade Center
159-1, Samsung-Dong
Kangnam-Ku, Seoul, Korea, 135-729
822.551.4380

IDC Market Research (M) Sdn Bhd
Suite 13-03, Level 13
Menara HLA
3, Jalan Kia Peng
50450 Kuala Lumpur, Malaysia
60.3.2163.3715

IDC New Zealand
Level 7, 246 Queen Street
Auckland, New Zealand
64.9.309.8252

IDC Philippines
703-705 SEDCCO I Bldg.
120 Rada cor. Legaspi Streets
Legaspi Village, Makati City
Philippines 1200
632. 867.2288

IDC Taiwan Ltd.
10F, 31 Jen-Ai Road, Sec. 4
Taipei 106
Taiwan, R.O.C.
886.2.2731.7288

IDC Thailand
27 AR building
Soi Charoen Nakorn 14,
Charoen Nakorn Rd., Klongtonsai
Klongsan, Bangkok 10600
Thailand
66.02.439.4591.2

IDC Vietnam
Saigon Trade Centre
37 Ton Duc Thang Street
Unit 1606, District-1
Hochiminh City, Vietnam
84.8.910.1233; 5

LATIN AMERICA

IDC Latin America
Regional Headquarters
8200 NW 41 Street, Suite 300
Miami, FL 33166
305.267.2616

IDC Argentina
Trends Consulting
Rivadavia 413, Piso 4, Oficina 6
C1002AAC, Buenos Aires, Argentina
54.11.4343.8899

IDC Brazil
Alameda Ribeirao Preto, 130
Conjunto 41
Sao Paulo, SP CEP: 01331-000 Brazil
55.11. 3371.0000

International Data Corp. Chile
Luis Thayer Ojeda 166 Piso 13
Providencia
Santiago, 9, Chile
56.2.334.1826

IDC Colombia
Carerra 40 105A-12
Bogota, Colombia
571.533.2326

IDC Mexico
Select-IDC
Av. Nuevo Leon No. 54 Desp. 501
Col. Hipodromo Condesa
C.P. 06100, Mexico
525.256.1426

IDC Venezuela
Calle Guaicaipuro
Torre Alianza, 6 Piso, 6D
El Rosal
Caracas, Venezuela
58.2.951.1109

IDC is the foremost global market intelligence and advisory firm helping clients gain insight into technology and ebusiness trends to develop sound business strategies. Using a combination of rigorous primary research, in-depth analysis, and client interaction, IDC forecasts worldwide markets and trends to deliver dependable service and client advice. More than 700 analysts in 43 countries provide global research with local content. IDC's customers comprise the world's leading IT suppliers, IT organizations, ebusiness companies and the financial community. Additional information can be found at www.idc.com.

IDC is a division of IDG, the world's leading IT media, research and exposition company.

02-144SYSTEM3386
July 2002

IDC
Analyze the Future

www.idc.com