White Paper

Ensuring Availability of Business Critical Applications

Introduction

Business data is the lifeblood of your company and assuring it is always available is the most important mission of your IT organization. System failures, human error, power outages and natural disasters put business critical information at risk of being inaccessible for extended periods of time or worse, permanently lost. You need solid protection against crippling data loss and application unavailability.

Continuous availability is no longer an ideal it is a necessity. Longer work days, expansion into new markets and customer demand for more efficient services create an expanded requirement for increased system availability. Users are demanding a means of ensuring very high availability of their applications and of access to data that permits them to accomplish their tasks and provide the highest levels of customer service. Interruption of workflow due to system failure is expensive and it can cause the loss of business. The need to increase computer system availability is becoming a key businesses concern.

The Cost of Downtime

Today's businesses and customers require high-availability solutions across the board and at an affordable price. A global business needs 24-hour access to information 365 days a year. In an Internet service model, organizations must anticipate customers arriving at their Web site and business partners interacting with their systems at any hour of any day, For many businesses, "regular business hours" have no meaning.The cost of downtime, whether unplanned or scheduled, can have substantial negative revenue impact both in terms of immediately lost business and productivity, as well as the subsequent effect of a potential loss of customer loyalty and confidence.

Business Operation	Average Cost per Hour of Downtime
Communications: Converged Services	> $10.0 million
Financial: Brokerage Operations	$6.45 million
Financial: Credit Card/Sales Authorization	$2.6 million
Media: Pay Per View	$150,000
Retail: Merchandise Sales	$140,000
Transportation: Airline Ticketing	$89,500
Media: Event Ticket Sales	$69,000

Source: Gartner, Dataquest, Contingency Planning Research and Others

What is the total cost to your business of one hour of critical system downtime? Can you afford this loss?

High Availability Clustering as a Solution

There is a solution. By combining standard servers running Linux or Windows with High Availability (HA) clustering solutions businesses can achieve between 99.99% and 99.999% uptime for business critical applications at a fraction of the cost historically associated with proprietary RISC-based systems. Thus, you can plan on between just 8 and 55 minutes of downtime, for both planned and unplanned outages, for an entire year. And this is for everything - from your mail server to your business critical financial management or manufacturing systems.

There are a number of HA solutions available today. The process of evaluating them and choosing one that meets your business objectives is simplified if you know a few key points about high availability and have correctly identified your business availability needs.

What Level of Availability Is Right For My Business?

To understand the HA needs of your business, use the following questions to do a self-assessment:

- Among my applications and services, which can I least afford to have unavailable to users and customers? Which most need HA protection?
- For each of these, what is the actual monetary cost associated with each hour of downtime? What am I willing to spend to minimize this downtime cost?
- How long would it take me to get the most critical business functions operational following a disaster hitting my data center?
- Should I invest in a high availability solution that enables me to resume business operations in minutes rather than days following some catastrophic event?

With this data in hand, you can identify your availability needs and make more informed decisions regarding the investment your business requires in order to ensure the uptime that your customers and business partners demand. With these needs in hand, you can move onto selecting the correct High Availability Solution to meet those needs.

Selecting the Most Appropriate High Availability Solution: Key Decision Criteria

There are a number of criteria which should be considered when selecting a HA solution. These range from the total cost of the solution, to the ease with which you can configure and manage the cluster, to the specific restrictions placed on hardware and software which can be supported. This section touches briefly on a number of these checklist items and then goes into greater detail on a few of those deemed most important.

Twelve Checklist Items for Selecting a High Availability Solution

- *Standard versions of Operating System and Application support*

If the solution that you deploy requires Enterprise or Advanced versions of OS, database or application software, the cost benefits of moving to a commodity server environment can be greatly lessened. With the proper HA middleware deployed, standard versions of applications and operating systems can be made highly available and can meet the uptime requirements of your business environment.

- *Support for a variety of data storage configurations*

In order to deploy an HA cluster, the data required by the application being protected must be available to all systems which may need to bring the application in service. There are a number of ways to share this data: via data replication, using shared SCSI or Fiber Channel storage, or using Network Attached Storage device. Whichever method you decide to deploy, the High Availability product used must be able to support all data configurations so that you have the flexibility to change as your business needs dictate

- *Ability to use heterogeneous solution components*

Some HA clustering solutions require that every system within the cluster be identical in configuration. This is common among hardware-specific solutions where clustering technology is delivered as a means to differentiate either servers or storage and among operating system vendors as they limit configurations they are required to support. This restriction limits your ability to deploy scaled-down servers as temporary back-up nodes, and perhaps to reuse existing hardware in your cluster deployment. It may be that deploying identically configured servers is the correct choice for your needs, but this should not be dictated by the HA solution provider.

- *Support for more than 2 nodes within a cluster*

The number of nodes which can be supported in a cluster is an important measure of scalability. Entry-level HA solutions limit you to a single 2 node cluster, typically in active-passive mode. While this does provide a level of increased availability by the addition of a stand-by server, it can still leave you exposed to application downtime. In a 2-node cluster configuration, if one server is down for any reason then the single remaining server becomes a single-point-of-failure. By deploying three or more nodes clustered together, you gain the ability to not only provide higher levels of protection, but you can also build configurations that are highly scalable.

Later in this paper, we will look at benefits available by building clusters of greater than two nodes.

- *Support for Active/Active and Active/Standby configurations*

In an active/standby configuration, one server sits idle waiting to take over the workload of its cluster member. This has the obvious disadvantage of underutilizing your compute resource investment. In order to get the most benefit from your IT expenditure, ensure that cluster nodes can run in an active/active configuration.

- *Detection of problem(s) at node and individual service level*

All high availability software products can detect problems with cluster server functionality. This is done by sending heartbeat signals between servers within the cluster and initiating a recovery should the signals cease to be delivered by a cluster member. There are another class of problems, however, which can be detected by advanced HA solutions. This set of problems include those which result from individual processes or services encountering problems which render them unavailable, yet do not cause servers to stop sending or responding to heartbeat signals. Given that the primary function of HA software is to ensure applications are available to end-users, detecting and recovering from these service level interruptions is a critical feature. Ensure that the HA solution you deploy can detect both node-level and individual service-level problems.

- *Recovery in-node and across-node*

As equally important as the detection of problems at both the server and the individual service level is the ability to perform recovery actions both across cluster nodes and within a single node. In across-node recovery, one node takes over the complete domain of responsibility for another. When systems level heartbeats are missed, it is assumed that the server which should have sent them is no longer in operation and other cluster members begin recovery operations. With in-node or local recovery, failed system services first attempt to be restored within the server on which they are running. Typically this will be done by stopping and restarting the service and any dependent system resources. This is a much faster recovery method which results in minimized downtime.

Later in this paper, we will look in more detail at the benefits of Local Detection and Recovery.

- *Transparency to client connections of server-side recovery*

Server-side recovery of an application or even of an entire node should be transparent to client-side users. Through the use of virtualized IP addresses and/or server names, the mapping of virtual compute resources onto physical cluster entities during recovery, and automatic updating of network routing tables, no changes need be required to client systems in order for them to access recovered applications and data. Solutions that require manual client-side configuration changes in order to access recovered applications greatly increase the time to recover and introduce the potential for injection of additional errors due to required human interaction. Recovery should be automated on both the servers and clients.

- *Protection for planned and unplanned downtime*

In addition to providing protection against unplanned service outages, the high availability solution you deploy should be usable as an administration tool to lessen downtime due to maintenance activities. By providing a facility to allow on-demand movement of applications between cluster members, you can migrate applications and users onto a second server while performing maintenance on the first. The need for "maintenance windows" when IT resources are unavailable to end-users goes away. Ensure that the HA solution you deploy provides a simple and secure method for you to perform manual (on-demand) movement of applications and needed resources among cluster nodes.

- *Off-the-shelf protection for common business functions*

Every HA solution that you evaluate should include tested and supported agents or modules designed to monitor the health of specific system resources: file systems, IP addresses, databases, applications, etc. These are often called "recovery modules". By deploying vendor-supplied modules, you benefit from both the run-time already done both by the vendor and by other customers and you have the assurance of ongoing support and maintenance of these solution components.

- *Ability to easily incorporate protection for custom business applications*

There will likely be applications, perhaps custom to your corporation, that you want to protect for which there are no vendor-supplied recovery modules. It is important, therefore, that there be a method for you to easily incorporate your business application into your HA solution's protection schema. You should be able to do this without having to modify your application, and especially without having to embed any vendor-specific APIs. A Software Developer's Kit that provides examples and a step-by-step process for protecting your application should be available along with vendor-supplied support services to assist as needed.

- *Ease of cluster deployment and management*

There is a common myth surrounding HA clusters: that they are costly and complex to deploy and administer. This does not have to be true. Cluster administration interfaces should be wizard-driven to assist with initial cluster configuration, should include auto-discovery of new elements as they are added to the cluster, and should allow for at-a-glance status monitoring of the entire cluster. Also, any cluster metadata must be stored in a highly available fashion, not on a single quorum disk within the cluster where if it should get corrupt or experience an outage the entire cluster will fall apart.

As shown above, there are a number of criteria against which you should measure any high availability solution. By understanding your HA needs and the items on the HA checklist, you can make the best decision for your particular needs. In the sections that follow, we will examine in more detail a few of the more important items for you to understand and consider.

The Availability Equation

The Availability Equation, TRESTORE = TDETECT + TRECOVER, illustrates how the total time required to restore an application to usability is equal to the time it takes to detect that an application is experiencing a problem plus the time needed to perform some recovery action. This equation introduces the key concepts of High Availability clustering, problem detection and subsequent recovery. In essence, HA solutions monitor the health of business application components and upon detection of problems, take actions to restore them to service.

Because the objective of deploying an HA solution is to minimize downtime, working to reduce detection and recovery times is key among the tasks of the solution that you choose to deploy. Since today's applications are in fact combinations of multiple technologies utilizing servers, storage, network infrastructure, etc., as you survey the options available, be certain to understand the technology used to detect and recover from all outage types. Each of these has a direct impact on service restoration times.

One technology critical to providing fastest possible restoration time is the ability to perform what is termed "local detection and recovery".

Local Detection and Recovery

In our earlier review of the Twelve Checklist Items, we touched briefly on the importance of service-level problem detection and recovery, also called Local Detection and Recovery. In a basic clustering solution, a number of servers are connected together and configured such that one or more servers can take over the operations of another in the event of a server failure. The server nodes in the cluster continuously send small data packets, often called heartbeat signals, to each other to indicate "I'm Alive". In simple clustered environments when one server stops generating heartbeats, other cluster members assume that this server is down and begin the process of taking over responsibility for that server's domain of operation. This approach is adequate for detecting failure at the gross, server-level. However, in the case of problems that do not cause the interruption or cessation of heartbeat signals, server-level detection is not just inadequate, it can actually magnify the extent and impact of an outage . For example, if Apache processes hang, the server may still send heartbeats even though the web server subsystem has ceased to perform the system's primary function. Rather than restart just the Apache subsystem on the same or a different server, as basic server-level clustering solution would restart the entire software stack of the 'failed' server on a back-up server. Thereby causing interruption to other users of the server, as well as extending the time to recovery.

Advanced clustering solutions provide a broad range of additional mechanisms to detect problems at a more granular level and enable recovery actions to be tailored to specific problems. Using local detection and recovery, advanced clustering solutions deploy health monitoring agents within the individual cluster servers to monitor individual system components such as a file system, a database, user-level application, IP address, etc. These agents use heuristics specific to the system component being monitored so that they can predict and detect operational issues and then take whatever recovery action is most appropriate. Often, the most efficient recovery method is to stop and restart the problem subsystem on the same server. This is much faster and less impactful, and therefore a less costly recovery method than migrating all application components to a stand-by server. By detecting failures at a more granular level than simple server-level heartbeats and by enabling recovery within the same physical server, the Time to Restore an application to user availability is greatly reduced. Make certain that whatever HA Solution you deploy can support local detection and recovery.

Scalability and Flexibility

There are a number of additional criteria against which you should measure any HA solution. Key among these is the ability of the solution to scale and flex as your business needs change.

Within the world of HA clustering there are many dimensions of scalability and flexibility.

For example, two methods exist for making your business data available to all nodes in an HA cluster: shared storage and data replication. Each is appropriate for certain environments, but how do you decide which is right for you? As important as deciding this is ensuring that whatever solution you choose can support both shared storage and replication configurations. So that regardless how your business needs evolve, the most appropriate storage configuration can always be deployed and the option to make any application highly available remains always open.

All high availability solutions support shared storage configurations where the data needed by clustered applications resides on either directly attached SCSI or Fiber Attached storage devices. All nodes in the cluster can access the storage device and as the cluster software migrates applications between servers during recovery, data paths between servers and the storage device are automatically reconfigured. While a shared storage configuration may be the correct choice for a number of deployments, there are some considerations that will make it less than ideal for others.

Primary among these is the expense of purchasing, deploying and maintaining a shared storage configuration. With prices ranging from several thousand to several hundred thousand dollars, small to medium businesses, as well as branch and departmental organizations may find shared storage to be cost prohibitive.

An alternative is to make use of data replication technology so that any data required by the application being protected is stored in the internal disks of all systems within the cluster. All writes of data to the primary disk of the system actively running the application are replicated to the disks of systems acting as stand-by nodes. This way, if the stand-by node should need to take over for the primary system, all data required is available locally. The cost and complexity of shared storage is eliminated while still allowing for full high availability.

In addition to cost, another factor in deciding whether to deploy a shared storage or a data replication cluster is the geographical distance between the nodes. Both SCSI and Fiber have distance limitations that for practical purposes limit the placement of nodes within the cluster to a LAN environment. If you are deploying a cluster between two locations across a WAN to provide disaster recovery protection for one of your sites, then you will make use of data replication to mirror your business data between the two facilities.

So, data replication technology plays two roles: it enables you to deploy an HA cluster at a lower price point and it supports the building of a stretch cluster across geographies for disaster recovery purposes. Be certain that the HA clustering solution you choose supports data replication clusters as well as conventional shared storage. With both data replication and shared storage configurations available, you can decide where among the clustering schemas shown in the Figure below you should enter the world of high-availability. By choosing a solution that can accommodate all deployment scenarios, you ensure that as your needs change, you can deploy the best solution for your environment.

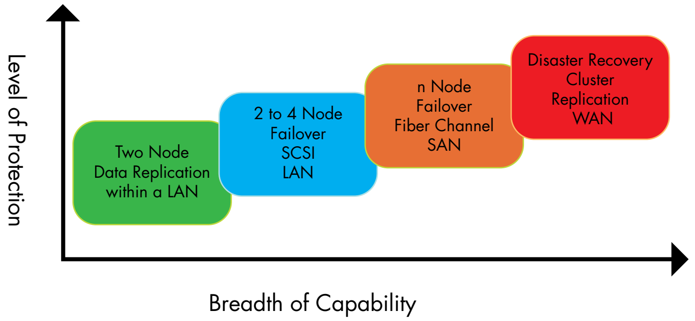
Level of Protection
Breadth of Capability
Two Node Data Replication within a LAN
2 to 4 Node Failover SCSI LAN
n Node Failover Fiber Channel SAN
Disaster Recovery Cluster Replication WAN

Greater than 2 node Support

Another scalability factor to be considered is the number of nodes which can be supported in a cluster. Entry-level HA solutions limit you to a single 2 node cluster, typically in active-passive mode. While this does provide a level of increased availability by the addition of a stand-by server, it can still leave you exposed to application downtime.

In a 2-node cluster configuration, if one server is down for any reason then the single remaining server becomes a single-point-of-failure. However, by deploying three or more nodes clustered together, you gain the ability to not only provide higher levels of protection, but you can also build configurations that are highly scalable. Two examples of such cluster configurations are commonly referred to as Many-to-One and One-to-Many.

In a Many-to-One configuration, a single stand-by server backs-up a number of active servers. If any one of the active servers should fail, the stand-by node will take over its operation. However, if a second active node should fail, that server's workload would then also become the responsibility of the single back-up server, thereby requiring it to be able to work in an active-active mode.

In a One-to-Many configuration, the domain of responsibility of the failed server is divided among a number of other servers in the cluster. In addition to enhancing overall availability by distributing points of failure, this approach of splitting a primary server's responsibilities in combination with active-active configuration offers significant economic benefits that further marginalize the cost of high availability.

This is because the ability to failover individual services in active-active mode to multiple servers eliminates the requirement for a potentially large, and therefore more costly server to be provisioned purely for the purposes of providing back-up in the event of failure.

Summary

Ensuring access to your critical applications and data is an imperative. Loss of productivity and revenue resulting from IT outages can cripple a business. Likewise, deploying critical applications on commodity servers makes tremendous economic sense. When coupled with the deployment of High Availability clustering software, you can realize the significant cost savings of a commodity environment without sacrificing application uptime. This establishes an economic model that supports the case for high availability solutions across a much broader range of business systems than could be justified for proprietary technical environments.

There are a number of criteria against which you can evaluate your choices for high-availability clustering. In making the decision as to which solution to deploy, the first step is to understand your business HA needs. Once your current requirements are determined it is vital to understand that these requirements will most likely change over time. Your selection criteria therefore will ideally incorporate the ability to evolve and make changes without requiring technology reinvestment or incurring disruption to users and customers caused by the need to reengineer to your environment.

A basic two-node, single vendor solution may meet your needs today, but will not provide investment protection as your business needs change in the future. Given the reality of "you get what you pay for" it is better to explore the full range of technical possibilities and make your selection based upon meeting current requirements today while providing flexibility for the future.

About SteelEye Technology

SteelEye is the leading provider of data and application availability management solutions for business continuity and disaster recovery for Linux and Windows and virtual environments.

The SteelEye family of data replication, high availability clustering and disaster recovery solutions are priced and architected to enable enterprises of all sizes to ensure continuous availability of business-critical applications, servers and data.

To complement its software solutions, SteelEye also provides a full range of high availability consulting and professional services to assist organizations with the assessment, design and implementation of solutions for ensuring High Availability within their environments.

SteelEye is a subsidiary of SIOS Technology, Inc. To contact SteelEye,
visit www.steeleye.com or call:
US/Canada 866.318.0108 Europe + 44 (0)1223 208701 Int'l + 1.650.843.0655

SteelEye Technology, Inc. 4400 Bohannon Drive, Suite 110, Menlo Park, CA 94025

SteelEye
TECHNOLOGY INC
www.steeleye.com

Replicate Any Data. **Protect** Any Application