1013

SharePlex: What is it, how does it work, and what are the benefits?

John Saylor
Quest Software, Inc.
610 Newport Center Drive
Suite 1400
Newport Beach, California 92660
(714) 720-1434

As dependencies on systems, applications and data increases, the need for a high availability solution increases. Today, many customers are operating in a 24x7x365 environment where systems, applications and data availability are critical to their business. Disaster tolerance and recovery planning are a necessity.

Hewlett-Packard's clustering implementations were rated one of the industries best by the Aberdeen Group, in December of 1992. The report, entitled "Clustering: an Alternative Growth and Operations Path," believes clustering is likely to be the foundation for commercial computing advances in the 1990's.

This has held true now through 1996. This maturing cluster technology has provided to be the basis for the much anticipated advances in enterprise-wide distributed computing. Quest Software has lead the way providing: horizontal growth, higher availability, and single-system views to users, applications and operators whom all benefited from their solution.

As long as business vulnerability is the key issue in technology-dependent environments without a high-availability, disaster tolerant solution in today's turbulent, confusing, and increasingly splintered information industry, executives can't make the right decisions based upon unreliable and incomplete information.

For executives, their most important strategic issue is to identify those key leverage opportunities that will deliver a major competitive advantage. SharePlex is a that strategic solution for achieving disaster tolerant high-availability systems in their competitive framework.

Executive Summary

In today's technology-dependent environments, information is often an enterprise's most valued commodity, providing a key opportunity to seize and sustain a competitive advantage.

So when a system goes down, the impact can be devastating. Many computing environments today are required to run 365 days a year, 24 hours a day in uninterrupted, continuous operation. If disaster struck, whether you currently have an underutilized system or a single data center environment, could your business recover?

Most businesses NEVER recover!

The Aberdeen Group has found that 50% of the companies dependent upon technology, never recover within 10 working days following a significant disruption.

The U.S. Government reported, 93% of the companies fail within five years of a data center disaster, that's a survival rate of only 7 out of 100!

In a September 1992 study, 300 U.S. businesses evaluated their level of vulnerability on a scale of 1 to 100; 100 being most vulnerable. The average rating for technology-dependent work environments reflected a vulnerability index of 75. In today's competitive environment, businesses cannot afford to be non-functional. The organization must meet business objectives, satisfy customer obligations, and continually maintain market share.

The solution to this dilemma is to provide capabilities unobtainable within a single system or data center. New capabilities must allow existing MPE or UNIX computer systems to work cooperatively and transparently with other systems, new or existing. These systems must appear to the user community as a single system and have the performance capability greater than a mainframe.

Enabling organizations to implement worry-free, business critical solutions cost effectively and to distribute their applications across multiple systems is a key competitive advantage.

First you must understand the problems...

A technology-dependent environment is one in which an enterprise relies upon current information to make informed decisions and depends upon the practical application of computer systems to perform normal business functions.

Disruption to normal business activities can come in many forms: natural threats, technical threats, and human threats. Most businesses are vulnerable; ask yourself some of the questions below:

What can happen (threat occurrence)?
How bad will it be if it happens (consequences)?
How often will it happen (frequency)?
How certain are the answers, to these questions?
Uncertainty . . .

When the workplace is crippled due to any disruption, in short, the business comes to a halt. This usually translates into disgruntled customers, lost revenues, and harm to an enterprise's goodwill.

We are surrounded everyday by enterprises that will be severely effected if disruption to their normal business activities occurs. Challenges faced by these enterprises must be dealt with to ensure they can provide service, meet schedules, ensure delivery and protect their share of the marketplace from competitors.

The Challenges a Single Data Center Faces...

A single data center or system is vulnerable, and at some point will impact your business operations. The effects on the end-user are key. By now, manual functions performed in the past are long forgotten. The business information entered in the computer systems is now what is relied upon to make all informed decisions.

Below are a few examples of failures or outages that could occur under normal operations.

System Interruption *15 minutes to 2 Hours*

- Operating System Failure
- Application Abort
- Network Interruption
- Power Outage
- Daily Backup

Long Term Outage *2 Hours to 8 Hours*

- Disk Failure RELOAD
- Hardware Failure (single point)
- Upgrades (OS, Hardware, etc.)

Disaster *8 Hours to 48 Hours*

- Hardware Damage
- Data Center Fire
- Natural Disaster
- Data Center Move

Executives today are looking for key opportunities that will deliver strategic information to decision maker's fingertips in a timely fashion.

Fault resilience is a critical element in the success of a technology-dependent data center. Anything that lowers the data centers availability decreases your resiliency.

Some of the common limitations that exist in a single data center include:

- The whole data center is affected by a disaster.
- Data redundancy is limited to Disk Mirroring or SPU Switchover, but both copies of data are lost by a disruption to normal business activities.
- Recovery from offsite tapes or, a business resumption service is slow and, all of the days irreplaceable transactions are lost.
- Fault-Tolerant systems provide hardware redundancy, but at a high cost, and loss of the data center still remains an issue.
- Service-levels drop during peak periods resulting in sluggish performance and the inability to balance the workload, requiring expensive upgrades.

So, now ... What is a SharePlex?

SharePlex is a strategic solution on both MPE and UNIX platforms for achieving disaster tolerant high availability systems over a wide-area creating wide availability. Operating in a 24x7x365 environment where system, and application and data availability are critical to perform normal business functions and reduce the enterprise's vulnerability.

The key features are:

- Shadowing
- Network File Access (NFA)
- Statistics
- Output Management (NBSpool)
- AutoRPM

In all these cases serious consideration should be given to a SharePlex-NetBase solution to meet the service demanded by today's enterprise. In some enterprises Service Level Agreements (SLAs) are setup to guarantee specific uptime requirements of the organization's computer systems.

Product Summary

Quest's NetBase product-line is the core of the Hewlett-Packard's SharePlex-NetBase, it resides on both MPE and UNIX architectures.

SharePlex-NetBase Bundle

SharePlex-NetBase Bundle provides customers with high system, application and data availability, disaster tolerance, shared system resources, and a horizontal growth strategy. SharePlex-NetBase Bundle can make resources on any system on the network available to all users. Files, databases, printers, and programs can be transparently shared among users on the network configuration, regardless of geographic location. SharePlex-NetBase Bundle can also enable customers to load balance between systems in their network, gaining optimal utilization of every machine.

Mid-range to high-end customers on either MPE or UNIX systems no longer have a top end with the Hewlett-Packard product lines. With SharePlex-NetBase these 996-800 and T520-12 systems have unlimited potential through a clustered approach. Many existing customers have run into various limits with operating systems, table limits, locking and concurrent end-user limits. This has all but evaporated by implementing a configuration which front-ends these processors' loads by distributing the user applications to this front-end and maintaining the core business processing as a database-engine on these high-end systems.

SharePlex-NetBase Shadowing

SharePlex-NetBase Shadowing is an entry level product which provides the shadowing component of the full SharePlex-NetBase Bundle. Customers needing the benefit of complete system, application and data replication, but are not quite ready to take advantage of all the features of the full bundle will benefit strongly from the SharePlex-NetBase Shadowing product.

SharePlex-NetBase Shadowing guarantees there is an up-to-date backup of all transactions. This, essentially, protects the user from loss of data.

SharePlex-NetBase Configuration

An HP3000 SharePlex or an HP9000 SharePlex solution requires a configuration consisting of a pair of at least two like to like HP systems (3000s or 9000s), networked together and running SharePlex. One system, acts as the primary system where the "master" copies of the files, databases, or applications reside. The other system acts as the secondary system where the shadowed, or exact duplicate copies of the files, databases, or applications reside. A SharePlex configuration can contain multiple secondary licenses, but only one primary license. All support issues are funneled through the primary license location. An exception to this would be a corporation who wanted a U.S. focal point and an additional European or Far East focal point.

SharePlex transport layer exists on top of the standard TCP/IP layer. The ThinLan Link data communication card and the ThinLan Link software are the only requirement to connect the computer systems. This is all standard hardware on all new systems purchased. Since this high-speed transport layer exists the overhead between systems is minimal and therefore line

costs are reduced. Specific network line requirements will vary among customers, but there are cases were one million transactions were sent across a T1 line and the transport layer only consumed 5% of the network. SharePlex is a wide-area solution there are other cases were transactions are shadowed even through a satellite. Basically, SharePlex is supported over any supported Hewlett-Packard connection were the TCP/IP layer exists.

Shadowing

The Shadowing feature of SharePlex-NetBase Bundle, and SharePlex-NetBase Shadowing provides the ability to create and maintain multiple exact copies of files and databases on different machines in the network configuration. Shadowing offers complete system and application environment replication. This allows for fast disaster recovery, performance off-loading, concurrent backups, and 24-hour access to business critical data. Shadowing is achieved by having a primary copy of each file or database and one or more secondary or shadowed copies of the master. All updates to the files or databases are performed on the primary system, with updates then transported and posted to the secondary system. SharePlex guarantees that they are exact duplicates. SharePlex-NetBase traps all updates to files and databases to ensure that the update is performed to the primary copy first, and then sent to all the shadow nodes.

In the event of a disaster, a secondary or shadowed copy of the applications and files or databases can be accessed immediately. Users simple reconnect to the secondary system and resume normal processing. Typically, PC users have a second icon and DTC users that have Openview DTC Manager installed simple perform a CNTRL-K and selected the secondary system.

Performance balancing or off-loading can be achieved by directing all reporting and inquiry requests for one group of users to the shadowed copy of the file, while another group of users can still access the primary copy.

Backups can be performed on the secondary system(s), while the primary system is being accessed. Transactions meanwhile are queued on the secondary system in case of any disruptions on the primary computer system. Customers can use the secondary copy to recover immediately from any interruptions. Additionally, online-backup tools compete with production systems for CPU resources. With backups performed on the secondary system production systems have reduced exposure to outside competitive processes on the system. A system with multiple CPUs for instance would lose a processor to this dedicated process for the duration of the online-backup. Secondary systems with online-backup usually have abundant resources to support this activity with minimal disruption to the users executing reports or inquiring the data on this system.

If the posting of data to a secondary copy is interrupted, the Shadowing feature automatically queues all updates on the secondary system until posting can resume. Transactions are posted on the primary and then queued on the secondary system so that if the primary were to fail during this time, the secondary system could post the queued transactions and have the exact copy of the primary systems files at the time it failed. High availability to applications and data is thus achieved and maintained, since there is always an up-to-date copy of the necessary file or database available.

Additional shadowing features of SharePlex support selective shadowing, bi-directional shadowing, multi-cast shadowing, consolidated shadowing and user-exit shadowing. UNIX shadowing will support automatic resychronization of data structures. Supported file types for MPE are: MPE flat files, message files, CM KSAM, NM KSAM, TurboIMAGE, IMAGE/SQL and POSiX files. Supported file types for UNIX are: HFS files, JFS files, NFS files, Oracle RDBMS, Sybase RDBMS and Informix RDBMS.

Network File Access (NFA)

Network File Access (NFA) gives users and applications transparent access to data and programs on other HP 3000s, on the network. This provides end users with access to more data

and applications than just those available on their local system, without having to know on which system the data or application resides. NFA also enables system managers to spread data and applications across systems, thus achieving better load balancing and cost effective horizontal growth.

NFA maintains a centralized directory of all files and databases that are available to network users. System Managers can enter a single command to allow a file or database to be accessible to all users. Therefore users are not impacted by having to modify programs, files equations, UDCs, or create programmatic sessions. When a program accesses a file or database, NFA checks the central directory for the location of the file. All the CPU usage to run the application is on the computer system that initiated the file request. All I/O requests handling access to a file is by the computer where the file physically resides. This avoids additional overhead from virtual terminals' sessions initiated from the local computer. HP's RFA requires the use of NS/3000 Services that places additional overhead of up to 30% on top of the TCP/IP protocol. NFA utilizes a data cache algorithm to reduce network I/O overhead so access across the network is optimal and transparent to end users. Transactions are prefetched in to reduce network overhead an increase performance. Closely clustered systems with FDDI links currently offer the optimal performance. This can be up to 10 times faster than Hewlett-Packards RFA (Remote File Access). Proper file placement is essential to optimize performance, and Statistics aids in this optimization.

Statistics

SharePlex-NetBase Bundle also provides a mechanism for capturing statistics on types of file access, network overhead, file access response time, and system location. As with any collection tool overhead can be high, so the collection of statistics can be simply enabled or disabled as required. Standard reports are supplied with SharePlex-NetBase Bundle, which condense and present the statistics in a clear and concise format.

System Managers utilizing the reports can determine optimal file placement on the network as well as using the reports to aid in profiling file activity by applications or users. A Scenario Generator is provided which can be used for modeling changes to file placement and network speeds prior to making changes in the actual configuration.

Another component of Statistics is an application tracing utility that enables developers to trace, in detail, all file system calls performed by a given program. The information provided by the tracing utility are type of call and critical parameters, such as DBGET. This tool proves invaluable when profiling an application to determine the calls made to the file system by older, legacy applications.

Output Management (NBSpool)

This Print and Spooling Management tool (NBSpool) *enables users anywhere in the SharePlex configured network to print to any spooled printer on the network at anytime*. This feature is especially beneficial during end of month report printing cycles, when it seems everyone in the department is printing. Print jobs can easily be directed or redirected to an available printer on the network.

NBSpool maintains a directory that is used to identify and locate spooled devices within the network. This directory can be altered real-time with no downtime to configure additional printers on the network. NBSpool, located between the application and file system, can create the program output file on the computer where the desired printer is connected. With NBSpool, the
local computer is free from having to transfer the spool file across the network, reducing network and CPU overhead. This strategy also avoids excessive disk I/O's copying the file from the local computer to the computer connected to the printer, even before printing begins. If desired, a local copy of the spool file can be made as a backup copy by using the COPY feature of SharePlex-NetBase Bundle. In addition, NBSpool provides automatic bannering and report distribution. A separate copy of the spool file can be generated for each

name on the distribution list. Each spool file will have its own banner, identifying the name of the intended receiver, the name of the spool file, the creator, and a custom message.

An additional benefit of NBSpool is that you can also dedicate a specific system as a print engine and offload all spoolfile creation overhead of 5-7% on each system to a single dedicated system. Also customers running SharePlex-NetBase Bundle need not purchase an additional network spooler to support the use of their Hewlett-Packard network printers.

AutoRPM

AutoRPM is an interface to Remote Process Management (RPM), which is supported by Hewlett-Packard. It is the only component of SharePlex-NetBase Bundle that requires NS3000 to be on the system running AutoRPM. Thus customers choosing to utilize the AutoRPM feature of SharePlex-NetBase Bundle need to purchase NS Services. Under MPE/iX version 5.5, NS Services is included on the FOS installation. AutoRPM grants system administrators the ability to easily configure programs to run on only a given system on the network. For example, system administrators may want to use this feature for performance reasons where a program runs best on a specific system in the network. AutoRPM requires two entries in the SharePlex-NetBase Bundle central directory, a logon record and a program record. The logon record defines the remote session logon information required by SharePlex-NetBase Bundle to logon to the remote system. The program record maps the local program name to a remote program name.

Once configured, AutoRPM intercepts all RUN and CREATEPROCESS MPE/iX commands. The RUN or CREATEPROCESS is then forwarded to the system specified by the logon record in the central directory. Upon completion of the program on the remote system, control is returned to the initiating system. Users are never aware that the program just executed on a different system in the network.

Why is it so difficult to get control of your environment?

Too many people look at the old school of thought which is to upgrade, upgrade, and upgrade. This solution offers the benefits of high hardware costs, higher software costs, inflexible growth and the same problems as before with the same solution...UPGRADE.

Today's computing environments should be installed to seize and sustain a competitive advantage so information is dispersed, in a continuous and uninterrupted operation.

Each time you, as a decision maker, reach this point, you are at a crossroads. You have an opportunity to regain control of your environment. The system you bought two years ago was a good decision at the time. Don't give it up so quickly! For a 30-40% increase in horsepower and similar cuts in support costs, you can have more.

What typically occurs on a single system CPU?

Once a CPU reaches 100%, application service-levels begin to fall off and work done by the computer systems is not done in a timely fashion.

Figure 1 - Single CPU Environment

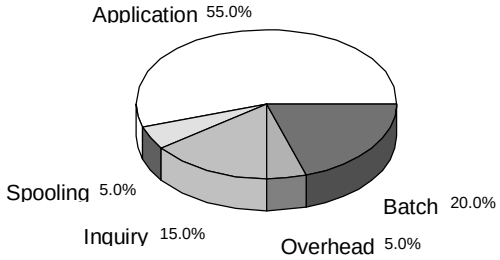

In this typically loaded environment, what would transpire if we were to introduce an Online-Backup? First, to create a sync point, minimal downtime would be a requirement to provide a rollback point. Once the sync point was established, the online-backup process would require an additional 30% more resources. The net effect to this environment would be an impact to all competing resource service-levels, resulting in slower response-times for

online and inquiry applications and an increase in the time to turn around batch reports.

To remain competitive, businesses are implementing an alternative approach. Enterprises are finding that maintaining copies of information throughout a network provides power and flexibility. Utilizing this approach, strategic business work groups are achieving significant savings and increased functionality in their environments.

How is this alternative approach implemented?

This approach is implemented through a balanced growth plan that focuses on the existing resource strengths and gears all new expenditures toward the price/performance of distributed clustered systems. An organization focused in this manner can sustain increased processing growth while keeping its new expenditures relatively flat.

When your system becomes overloaded! Consider adding another processor and implementing a network. Take the pieces of this "processing pie" and distribute a typical single CPU environment load over two or more systems. Create multiple CPUs independently controlled as symmetrical processors to provide independent redundancy. Redundancy that is not built into the same physical piece of equipment, which allows specialized servers, reduces software licensing costs, and provides the performance necessary for that work group.

This growth is horizontal, which makes use of past expenditures and minimizes future cash outlays. This allows organizations to smoothly transition between technologies. It also protects the business from quickly changing technologies until it fully understands the benefit to the business environment and the new risks associated.

After all, the main goal is to provide availability to enterprise-wide information, where the users should not need to know about:

- The access process or concurrence
- The location of the systems or files
- Technology of the systems
- Failure and recovery process
- Performance balancing

This helps the enterprise to reduce cost, increase information availability and provide a flexible, cost-effective solution. Some of the elements unobtainable within a single data center, utilizing redundancy are:

- Disaster tolerance
- Data replication over long distances, i.e. WAN
- No single point of failure
- Real-time data availability, 24 - Hours
- Online backup w/o disruption, 24x7 data center
- Second copy of data for OLTP applications is immediately accessible

Some of the elements of clustering multiple systems together to provide unlimited growth and balance are:

- Focuses on existing strengths
- New growth benefits from price/performance
- Clustering allows modular horizontal growth with specialized servers for adding new processing power or offloading systems
- Molding your computing environment around your business functions, not centering your business around your computers
- Software license costs are reduced, since you are sizing your processing power to the specific user requirements
- Leverage off legacy resources
- Integrate new technologies slowly

Is high availability important?

Many organizations may have started by connecting a few users together on a single local-area network, pretty simple stuff. As those organizations grew, they had to figure out a way to connect hundreds of users together. They needed to do this cost effectively and providing high enough service levels for maximum productivity. Even now, many of you may be facing this challenge.

Choices now must be made to determine the critical needs of the organization. Management must set priorities on processing and operation requirements and eliminate non-critical applications. You must focus on which daily business functions could you afford to lose without financial loss. Then determine the resources these functions require to restore adequate levels of service once an outage occurs.

Following a disaster, the business must focus on its resumption. A solution must be in place to bring the users back online in minutes. This greatly minimizes the loss of access to data, saves tens of thousands of dollars per minute, and eliminates long-term crippling effects on an enterprise's ability to do business. Most companies fail to adequately estimate the time required and effort involved to reassemble their work environment, resulting in a business impact.

How do you protect your technology-dependent environment?

In a technology-dependent working environment, a clustering solution survives a disaster. However, most system and data outages are not caused by disasters, but are attributed to hardware, software or networking failures. A well-implemented clustering solution should restore close to normal operations within one hour. This can give a business a distinctive competitive advantage by, reducing both planned and unplanned system outages.

Hewlett-Packard's recent consolidation of its own data centers was profiled in "Hewlett-Packard Builds Logical Mainframe with NetBase," on May 1993. It is heralded as one of the ten largest Information Technology Centers worldwide. Built on an industrial strength network, HP is integrating any number of MPE, UNIX and Desktop systems or networks. The ITC had to maintain a high-level of support and service that HP has always provided to its users. Around-the-clock availability was the HP requirement, supporting over 28,000 users in the Americas, and accessing more than 300 applications on over 150 systems.

The most noted application is the HP Response Center's ability to shadow critical call tracking information between the Atlanta, GA and Roseville, CA locations. The key was to maintain a high-level of support and service to the users of Hewlett-Packard equipment and systems.

In the profile, Wendy Odlum, an Information Technology Engineer observes, 'NetBase is an integral part of the infrastructure of the data center. It's invaluable and invisible - and it allows me to do more with less, exactly what everyone is striving for these days. NetBase allows our support organization to be better, faster and cheaper. It has proved itself as both a strategy and a tool.'

There are many changes in the working environment as HP moves towards a distributed computing environment. The reference center concept is proving itself day after day - saving disk space, money and people. As things change, HP has been able to absorb the changes very quickly with its reference machine strategy and NetBase. The company has not had to wait on new equipment or more people. 'NetBase makes the link between machines invisible. That's the whole point. It has allowed us to create a logical mainframe servicing our users, all of whom think they are accessing the information locally,' states Wendy.

How would your business be effected if disaster struck?

Your organization should perform a business impact analysis to determine the potential financial impact of a major disaster. Next, you must evaluate which solution makes sense to protect the business. Loss scenarios should be developed. All types of losses should be considered, accounting for every major item and how they effect availability. It is important to understand the MTBF - Mean Time Between Failures of your organization and more importantly how they compare with your competition.

The shaping of future outcomes is an important factor in determining the critical needs of the organization. Management must set priorities on processing and operational requirements. An application approach is usually best, eliminating non-critical applications first, then determining which applications must be recovered in parallel, considering all the end-user priorities.

These recovery strategies are key in determining the feasibility and speed with which the business can recover. The order of priority is usually as follows:

- Communications
- Computer Systems
- Personal Computers
- Facilities
- Work Environment

No one solution can solve all the computing challenges you may be facing, but as you see there

are several excellent choices available. Your business may very well depend on the choice you make today in a time of crisis.

Recently, L.A. Gear, an athletic shoe and apparel manufacturer in Santa Monica, California, survived a brush with a 6.6 earthquake which badly damaged the building where its data center resided. The county safety commission ordered L.A. Gear's building to be cleared and evacuated for several days. However, no orders or other valuable data were lost since L.A. Gear shadows its entire computing environment to another set of computer systems at its Ontario, California, site. This site was not effected by the quake. Without the disaster tolerant solution, the whole company's operations, including the corporate staff, could have been shut down for days, and the company would have experienced a considerable financial loss.

A remote or hardened data center is the current attraction for business recovery preparedness. Many customers fall short since the recovery from offsite tapes or, a resumption of service is slow and all of the day's irreplaceable transactions are lost. However, a few of these data centers are quickly adapting to changing trends. They are providing a full service resumption, where businesses can shadow a copy of their critical data into a remote or hardened data center. They then have the ability to bring their users back online in minutes with no loss of data. This type of business resumption provides businesses with a real solution.

The success of these approaches attains new levels of usability, allowing organizations to provide capabilities unobtainable within their current single data center environment. Challenges like HP's and L.A. Gear's are faced by many organizations where business recovery preparedness is quickly becoming a survival requirement.

This article, is contributed by John Saylor of Quest Software. John was previously, a specialist in the areas of Capacity Planning, Performance, High-Availability, Disaster Recovery, and Data Center Operations at Hewlett-Packard for over 11 years.

For additional information on how these types of products can assist you in your specific environment please contact John Saylor at (714) 720-1434 or send an e-mail message to jcsaylor@quests.com.

References:

1. *Business Vulnerability: In Today's Technology Dependent Environments (John Saylor - April 96, May 94)*
2. *HP 3000 SharePlex/iX - NetBase Data Sheet (5963-6617E)*
3. *The Aberdeen Group, in December of 1992. The report, entitled "Clustering: an Alternative Growth and Operations Path"*
4. "Hewlett-Packard Builds Logical Mainframe with NetBase," on May 1993
5. *All manuals and data sheets (Quest Software June 96)*