Presentation #: 1000
Title: High Availability and the HP3000

Ralph Bagen
Aircast Incorporated
92 River Road
Summit, NJ 07901
(800) 526-8785 x227

Perspective

More than any other time in history, MIS is having a larger impact on the business operations of a company. The industry is moving so fast, and the pace of change in so furious that the opportunities for a business to harness a new emerging technology can translate into a significant competitive advantage. Put another way, the cost of not staying abreast of technology can seriously jeopardize a company's ability to stay competitive.

One such method of using technology to its fullest is to ensure that the computer systems that your company have invested in are utilized optimally. This loosely equates to systems being made available longer and performing faster. It is the purpose of this paper to introduce some of the technologies available today for enabling these "new requirements for business".

Information Availability

In today's business environment, a large value is being placed on data being accurate and available. The longer that it takes to deliver say a large month end report directly affects the ability of a company executive to make a timely and informed business decision. Today's executives need strategic information on demand 24 hours a day. This requirement has led to an increase in the popularity of Data Warehousing technology. In addition to this, a common requirement of such technology is to search or "Data Mine" for patterns or trends that are not immediately obvious or even suspected.

To say the least, the cost of downtime is becoming a very costly proposition. Information Services are being required to have "high availability" - sometimes even 24 x 7 propositions. Before, such levels of service were only afforded to mission critical operations. Now, with reduced hardware prices, lower cost of ownership and the number of applications multiplying, fault tolerance is no longer just a consideration for large MIS environments. Such a business pressure must also ensure a guarantee for a fast and safe disaster recovery plan.

What is Fault Tolerance?

The definition I like to use is the following:

"The ability to maintain computer systems functions at all times, regardless of peripheral, componentry or even CPU failure - sometimes including complete redundancy to provide disaster resilience."

Limitations of a single computer

The load put on systems normally only increase over time. When this happens, sluggish performance leads to constant pressure for upgrades. This in turns raises questions over the competency of an MIS organization. It also represents a vicious cycle.

Even an environment that uses disc mirroring or some other level of disc redundancy (Levels of RAID) are vulnerable, since when an unplanned event / disaster occurs, both copies of the data are lost. Even when service resumption occurs, there is the possibility of transaction loss which causes a time consuming data verification effort just in order to identify a problem.

Also, when a single computer system gets overloaded (i.e. during peak periods such as month ends), service levels drop which results in a decrease in productivity and an increase in frustration for the work force.

A typical HP3000 environment is a classic victim of its own circumstance. An HP3000 excels at processing high transaction volumes while servicing users. When a CPU gets overloaded, system response time suffers since batch jobs and on-line users are competing for the same resources. Tools like Workload Manager (with MPE/iX 5.0) have helped enormously but it is still a balancing act that has to be maintained to keep everything flowing.

Generally what happens is that just before frustrations boil over, an upgrade proposal is made to management. I personally went through this process 3 times before figuring that there has to be a better way. The critical piece in the argument was that the classic environment has no disaster tolerance. Even something as essential as having good backup strategies in the event of a disaster was in direct conflict with the objective to keep on-line systems as available as possible. Traditional "disruptive backup methods" render a system essentially unusable during the taking of backups (see TurboStore/iX)!!!

Advantages of clustered systems

What is clustering? The Aberdeen Group's defines it as "a multi-node computer system which has a single view as seen by users, programmers, operators and administrators". A clustered system represents multiple CPUs configured in such a way that the entire cluster is not vulnerable to a single point of failure, thereby theoretically increasing system availability. This is accomplished by making multiple copies of the data available (immediately for OLTP environments) via data shadowing. By having multiple CPUs, it is often possible to actually reduce the cost of software, because of tiered software pricing or multiple CPU pricing schemes. Often, true on-line, non-disruptive backup utilities are available.

What are the alternatives?

One of the first options to consider when your system appears to be overloaded is to identify what processes or applications are causing resource constraints. Then it must be determined whether or not these "problems" can be re-written, re-scheduled or moved to a different environment / machine. One way to accelerate the speed of your existing system if you haven't already is to implement a network. The sheer speeds at which data can now travel will alleviate many "bottlenecks" that are felt in a classic HP3000 environment.

If there is no recourse but to upgrade, consider configuring two smaller CPUs into a cluster, rather than going for the larger CPU. As stated earlier, software costs will in all likelihood decrease. This is particularly true if you are presently sitting on platform dependent application tier boundaries. Perhaps more important, however, is the fact that if the applications are correctly straddled across the "cluster", each smaller CPU will actually outperform its more expensive counterpart. This is essentially accomplished by taking pieces of the "processing pie" and distributing them across the clustered system. One such solution is Netbase (SharePlex/iX).

Distribute batch processing

Using Netbase's data shadowing feature, we moved most batch processing to other processors in our distributed environment. This allowed two major advantages to occur: batch jobs no longer impede our mission-critical OLTP environment and we can run large CPU intensive jobs in the middle of the day / all day without users suffering in productivity. The latter is critical if you intend to have an environment performing some kind of decision support function.

By having a shadowed copy of your data on another processor, 24 hour availability is attainable since you have a built in backup of your data at any time. If you prefer, any copies of the data can be backed up at any time without users requiring sessions or jobs on the other systems having to log off. This will mean that no additional backup utilities are really needed. Obviously, multiple copies of the source data will serve as a backbone for a disaster recovery scheme, but a modular approach to application distribution may also serve to simplify the function of improved data / system security.

Distribute on-line inquiries

Since on-line inquiries can be performed on any copy of the data, the system has multiple options for returning the answer to a query. It can be routed to a particular machine that is optimized for inquiries or it could ascertain which machine in the cluster is the least busy and pass the request to it. Either way, your computer environment is being used in a more optimal fashion, probably at a cost that is less than currently being expended.

Such distributed / replicated environments are becoming increasingly popular. SharePlex is the method by which Hewlett Packard themselves have interconnected over 260 systems into a world-wide network to provide a seamless environment for handling calls to their response centers. Each time zone handling a call is working on a shadowed copy of the data from the response center where the call was originally logged. This provides for a lack of dependence on a single centralized computer system while at the same time allowing each call handler instant access to data without establishing a connection to the computer system where the call was originated.

Data Shadowing

Data shadowing permits concurrent copies of data and programs to reside and be used on multiple CPUs spanning multiple facilities. The data can be replicated over LANs and WANs. NetBase shadows many different types of files (IMAGE, KSAM, MPE files, AllBase/SQL & programs). Other products exist such as ORACLE Replicate for ORACLE 7.0 or higher. Most shadowing techniques can automatically attach users to the most effective copy of the data, permitting a fast recovery from data center outages. One particular extension is known as consolidated data shadowing. This method creates a read-only copy of enterprise-wide data which is completely transparent to the applications. This may be the simplest way for your organization to provide centralized inquiry and reporting functions while at the same time creating a central data source or warehouse populated from multiple, disparate databases.

Other HP High Availability Features

In the past several years, HP has recognized the emerging need for high availability of your company's production environment. Many tools and utilities have been incorporated into the fabric of the operating system itself. Other technologies have been made available via third-party providers. Let's examine some of these features:

Non-disruptive, on-line backup software (TurboStore/iX II or Backup/iX from Orbit Software) will allow for backups to occur with full production system activity occurring. This is accomplished basically by allowing for a sync point or data acquiesce that would allow for a rollback posture if needed.

IMAGE/SQL possesses a dynamic dataset expansion (DDX) facility that will allow for databases to be sized on the fly when the get to be within user defined acceptable thresholds. This utility is recognized by Adager and Quest Software's Netbase products.

Mirrored Disk/iX allows for disc drives to be paired and mirrored 100% - this is also known as RAID Level 1.

Utilities like setclock, Auto Restart/iX and faster boot times are currently available and soon (MPE/iX 5.5 & 6.0) Patch/iX will allow for a system to dial out and based on its configuration will bring in patches that are deemed appropriate and actually patch the system if available. Also coming are SCSI device "hot swaps" and on-line configuration of peripherals.

RAID - Redundant Array of Inexpensive Disks

RAID arrays can be configured in several different ways:
Level 0 (none): Block striping provides high performance by balancing disk activity across all disks.
Level 1 (disc mirroring): Drives are paired & mirrored 100%
Level 5 (block striping): Striped with parity information distributed equally across all disks for data redundancy.
Level 6 (block striping & mirroring): Every block has a mirrored block.

Components of Clustering

Network File Access (NFA) allows remote file access transparently to users.

NBSpool / Print Services is a router that provides seamless access to any printer in the cluster.

AutoRPM gives users transparent access and execution capabilities to any program or application in the cluster. Users are redirected transparently via NS/3000 to remote systems to execute applications. This allows anyone with security to run any software on the network. It has the benefits of not incurring cost of carrying multiple copies of the same application software, and in fact means that money can be saved by putting the software on the CPU with the lowest tiered software pricing. It also greatly simplifies the task of network application integration.

No plan, no business

Continuing to fulfill customer obligations and commit dates require that computer functions stay up. Just about every business function today is handled by at least one computer system in your organization (including your phone system or personal computers). Business recovery preparedness is quickly becoming a requirement for survival. It is critical that the business function is resumed as rapidly as possible after any down time / disaster. It has been shown that about one third of companies who lose computer services for more than 2 days are completely out of business within 1 year. It is also sometimes very difficult to quantify damages done in customer relations just by losing a vital computer function for an hour (what do you do when you get a busy signal at a mail order / catalog house?).

Create a disaster recovery plan

This plan must be formally described in order to be effective. In the plan, describe the actions to be taken in the event of a serious disruption of normal business activities. A team should be assembled that will determine what the critical needs of the organization are. This might be a view from an application based perspective. Based on this, priorities should be set for processing and operations. It is imperative to secure top level management commitment for the process of formulating the plan. As a team, the sources of risks should be assessed. For example, should you plan contingencies for natural (fire, flood, earthquake) or even human threats (sabotage, bombs, theft)? Remember that in many scenarios, employees may not be able to actually get to the organization. Are there threats that exist whereby your company's systems survive but others do not (i.e. phone company, power utilities)? The plan will represent what level of risk your company is willing to take - establish what "reasonable precaution" means in your organization. The cost of formulating and implementing your company's disaster plan should be less than the cost of your systems being down for a pre-defined period of time (there are exceptions to this generality, i.e. hospitals, defense systems).

This must be an iterative, evolutionary process to be successful. The most common mistake is to fail to test the plan. There must be a scheduled test of the plan, most effectively unannounced wherein every aspect of the plan is rehearsed, even if this means telling employees to report to a different city or stay in a hotel. Costs in testing disaster plans are large. Failure to test disaster plans are fatal. In order to be effective, I feel plans should be tested at least once a year, preferably every 6 months. The frequency for testing is different for each company but is related to the amount of change that occurs within it. Each functional change in your business will result in a revision and re-test of the disaster plan.

How Do I Get Started?

One way is to attended conferences like HPWorld. There are special interest groups (SIGs) that deal specifically with topics such as High Availability. HP prides itself in receiving feedback from its HP3000 install base. Figures 1 - 3 are part of a study that I participated in with the High Availability SIG, and it shows the input process to which HP has channeled some of their R&D and resource efforts in the area of HP3000 High Availability. Another method is to approach your HP Sales Representative or HP Channel Partner. Solutions are readily available, and the cost of building some entry-level redundancy into your HP3000 environment may very well be worth your time investment. Don't get caught up in a perpetual "how long can we get by until we have to upgrade" cycle!

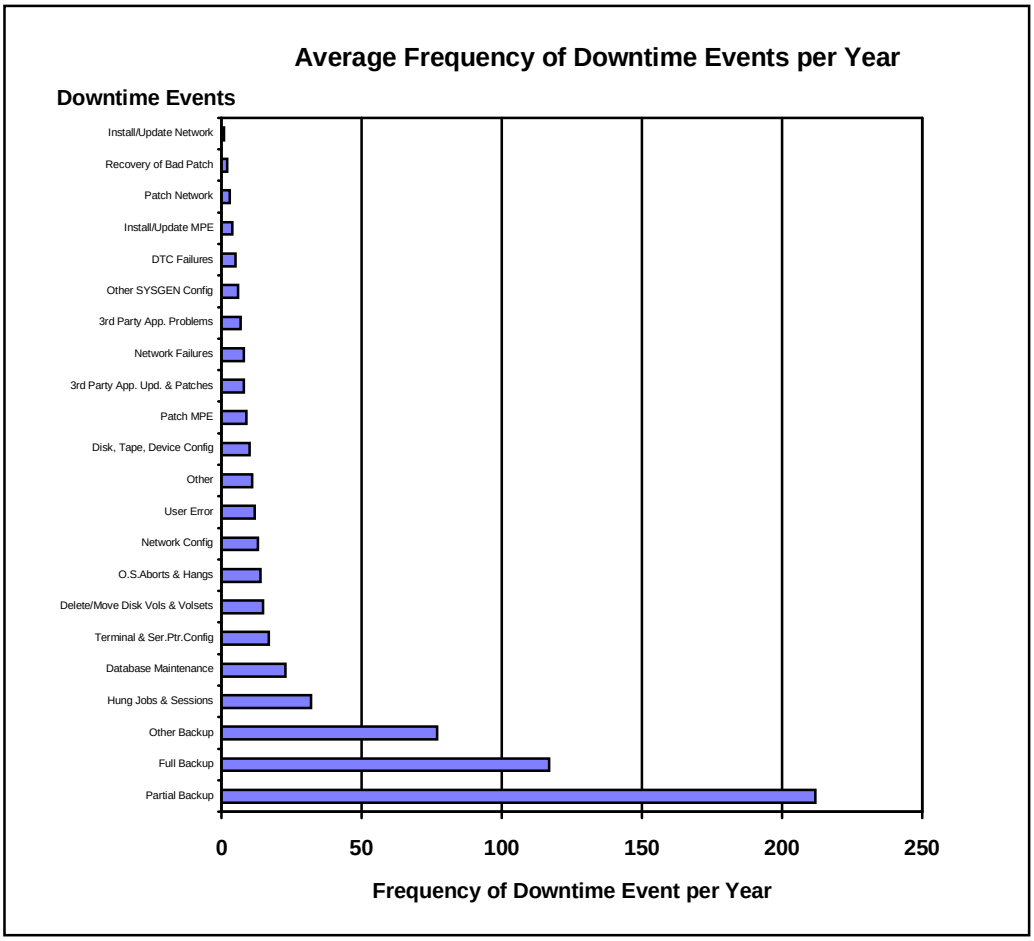

Figure 1- CSY Lab 7/94

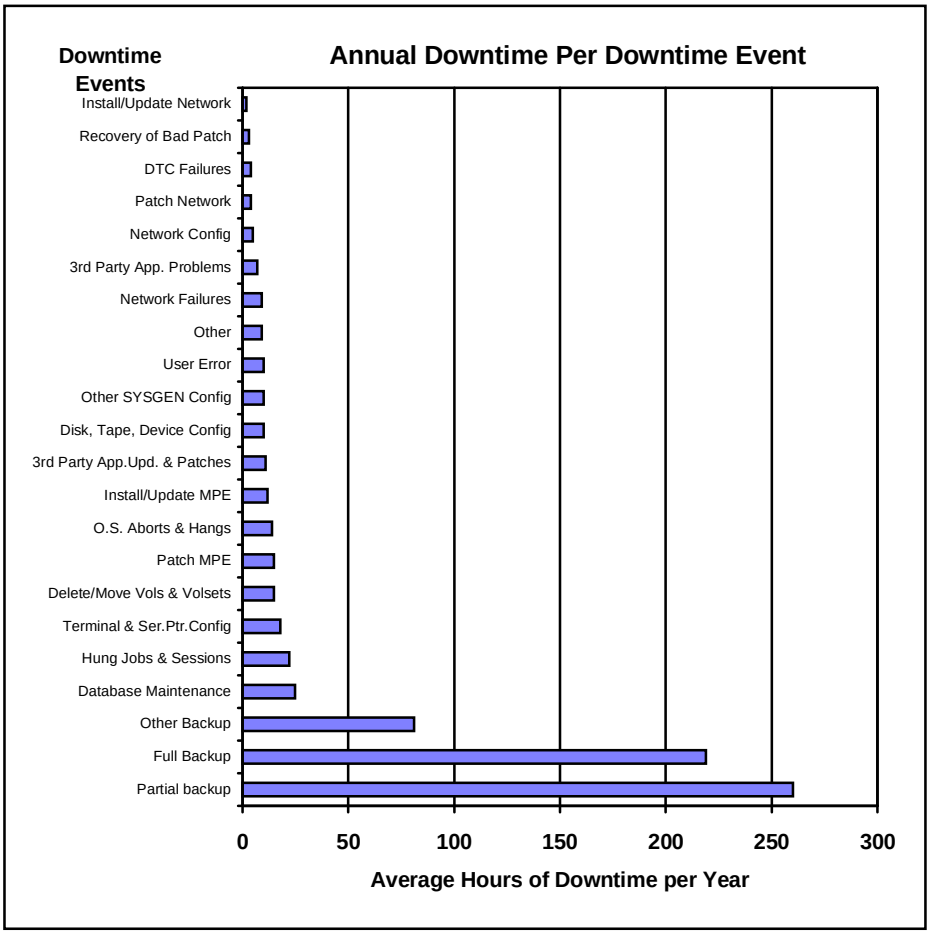

Figure 2 - CSY Lab 7/94

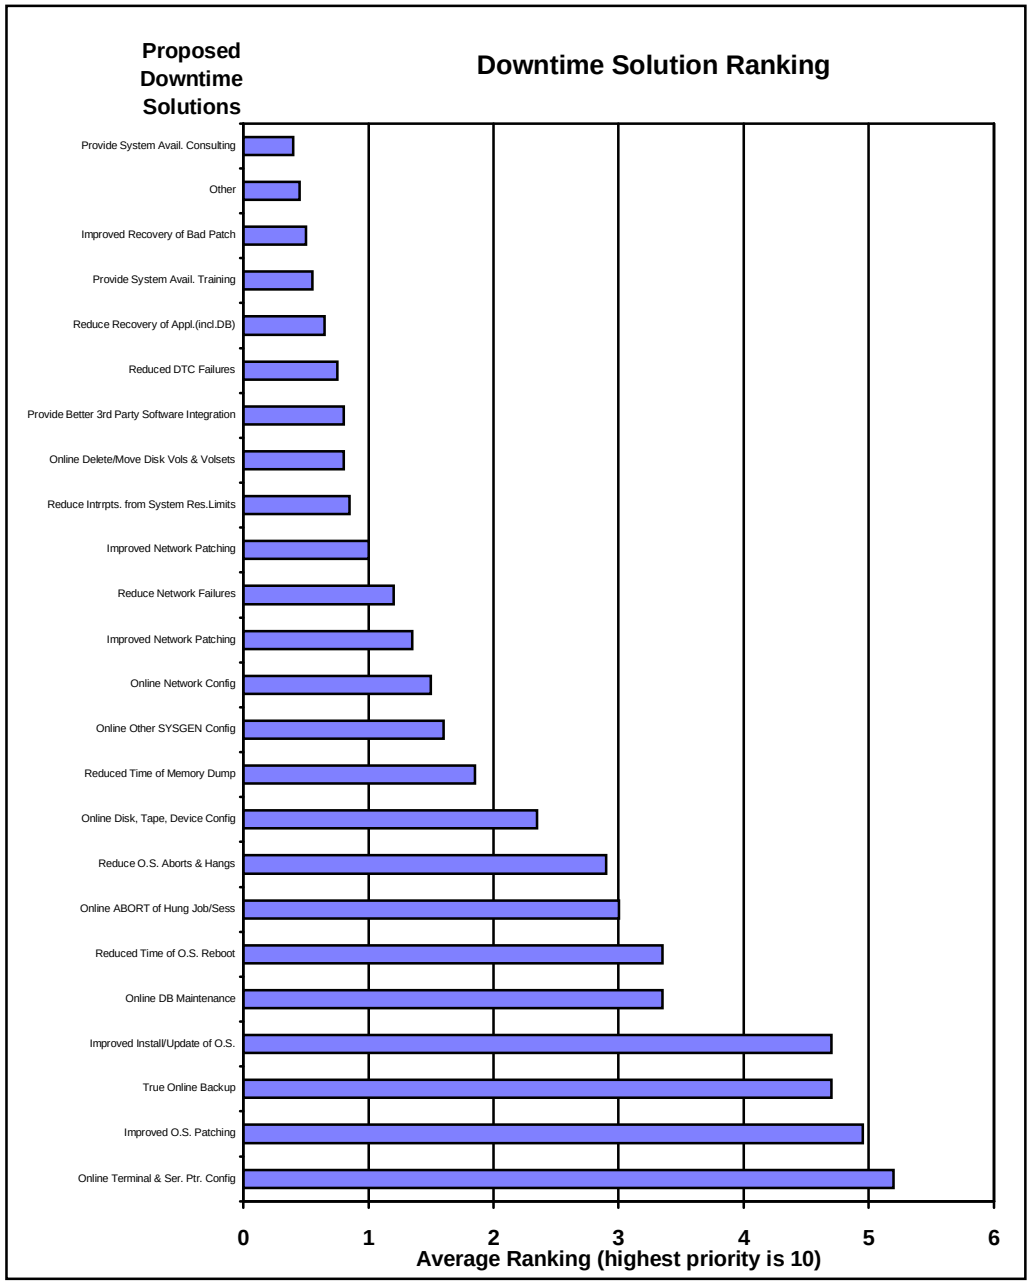

Figure 3 - CSY Lab 7/94

BIOGRAPHY FOR RALPH BAGEN

Ralph Bagen is the Systems Manager at Aircast Incorporated, a privately - held manufacturer of orthopaedic products located in Summit, New Jersey. He has held this position for eight years. where he has been responsible for the design and integration of computer based systems. Ralph has managed several projects such as the implementation of MRP II, installation and administration of a Novell 3.12 local area network, and has integrated the data communication requirements of the three Aircast manufacturing facilities. He is presently involved in the implementation of a client / server OLTP order entry system along with EDI and is evaluating computer telephony applications that would enable the HP3000 to talk to the telephone PBX via the PC clients. Ralph holds two Master of Science degrees, one in Mathematics from Fairleigh Dickinson University, and another in Computer Science from New Jersey Institute of Technology.