

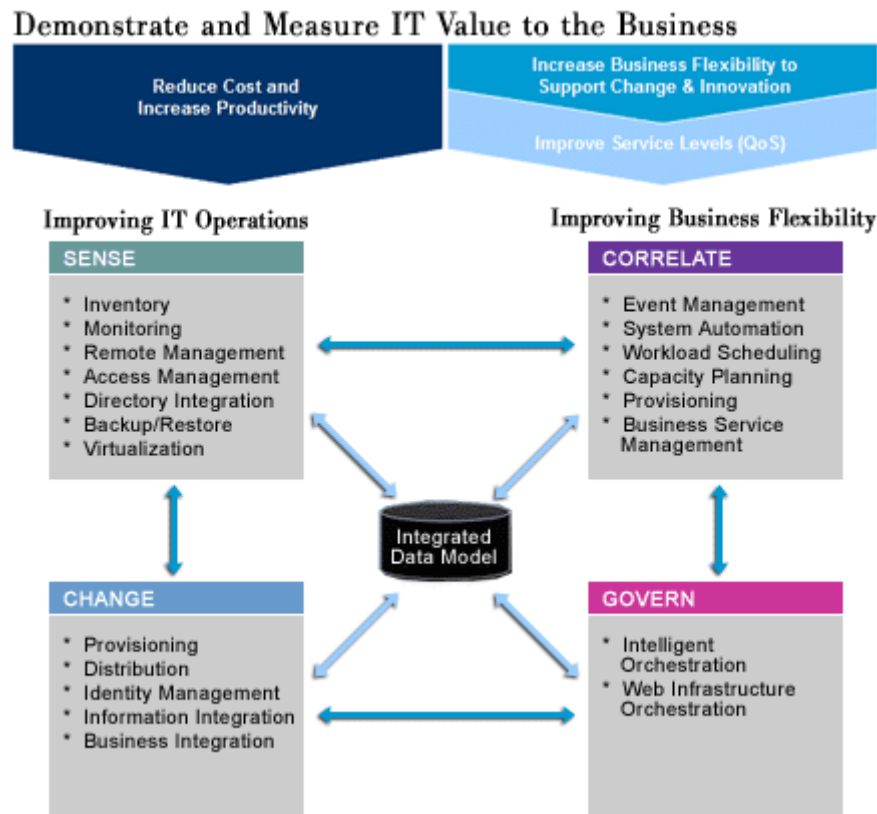


Does Highly Available also mean Under-Utilized?

by Diane Hage
Software Sales Manager
Datatrend Technologies

In a distributed environment, it's quite common for companies to purchase more than they need in order to support their individual highest user demands. This over-provisioning typically can yield an average server utilization of less than 20%.

In an on-demand world, your company needs the flexibility of allocating resources of servers and storage as they are needed to manage to the peak times and reallocate when the peak is past.



But what does it take to be that dynamic in today's on-demand world? Well, for starters, ask yourself, what are you really looking to accomplish? If what you are truly looking for in an AIX server environment is robust high availability and disaster recovery, with fast failover and streamlined clustering, then the HACMP, HACMP Extended Distance, or HACMP Smart Assist for Websphere might be your choice. High-availability also means less down time, especially in your mission-critical areas.

The First Step

The Standish Group International, a market research firm specializing in e-commerce, conducted a study to try to quantify the cost of data outages. They estimated that for a transaction processing system, the cost associated with an outage is approximately \$2,500/minute, or \$150K per hour at a minimum.¹

If your concern is setting up your storage to be highly flexible, the SAN Volume Controller (SVC) can help keep your storage costs manageable while flexible. If it is your Intel Environment which you're looking to virtualize, then VMware might be your solution.

Here are some other areas to think of when designing your highly available system. What about:

1. Building a system from test to production
2. Taking a standard desktop installation and automatically deploying and redeploying over new or newly assigned machines

3. The best practices skill set of your top people:

What if your system needed to be rebuilt and your Systems Administrator had left your firm? Would you have automated the workflow and processes that could just emulate what he/she does? Would the system be able to automatically ensure all steps are taken in the right order and consider all contingencies?

With IBM's Intelligent Orchestrator, modeling templates and customizable templates can allow any level administrator to work through such a crisis as those mentioned above.



How many servers do you need to accurately test your planned environment? IBM Provisioning Manager can enable production scale load testing, repurpose a small pool of servers for any number of application staging environments. Then they can be repurposed to load test on other environments or be put back into productive use where they're most needed.

How frequently do you need to load a fix pack or software throughout your system of distributed machines? With this provisioning software, it will automate that process, or even cluster a rolling upgrade on your Web server cluster and application servers.

In April of 2004, IBM announced the IBM Virtualization Engine (VE), which consists of Systems technologies and Systems Services, allowing:

- Multiple operating systems to be supported in one server
- Dynamic logical partitioning (dynamic LPAR) which allows server resources to be grouped into logically separate systems within the same physical footprint
- Virtual I/O, allowing partitions with a server to communicate via a high-speed virtualized network.
- Virtual LANs to enable high-speed communications between partitions
- Virtual disks in ESS and FASTT

Protecting Your Organization

IT organizations are being held accountable these days to be more efficient, more flexible, and more timely. All these issues mean a concern for automation of more complex systems with fewer people, of automation of processes to lessen the complexity of technological changes. So making your systems highly available can mean many things. One thing is sure: you can't afford for your systems to be down. While everyone knows the horror stories of web-centric companies whose on-line e-commerce systems crashed at peak customer buying time, the costs are equally as devastating for data warehousing systems: estimated by the Standish Group as approximately \$6,800/minute during peak loads when decision-support deliverables are being generated.¹

Datatrend can help you in the evaluation process, the infrastructure analysis, and the implementation of your choice of solution. Why? ... because we've done this time and again with many customers and understand how to be flexible and earn your trust in this important component of your on-demand business. Give us a call, and let us learn more about where you are at in your highly available business, and let us help you discover and implement the right design just for you!



¹ DB2 Universal Database and the Highly Available Data Store, Paul Zikopoulos, James Stittle, Roman Meinyk, 9 October 2003, IBM developerWorks.