

Planned Downtime: How Your Disaster Recovery Solution Can Reduce Costs and Shrink Outage Windows

By Sherri Atwood

TODAY many companies rely on clustering, volume management and replication technologies as a line of defense against unplanned downtime—server failures, site outages, and other natural events that threaten customer service levels. These technologies can also be leveraged to reduce the costs and outage windows associated with planned downtime events—providing a significant return on investment (ROI) bonus.

THE ESCALATION OF PLANNED DOWNTIME

Planned downtime refers to any scheduled administrative operation that could potentially disrupt service, such as system upgrades, the transfer of data, server consolidation and site maintenance. Planned downtime occurs far more often than unplanned downtime, partly because hardware systems are more robust and resilient than ever, and mean time between failures is constantly improving.

The frequency of scheduled downtime processes is definitely increasing. A large enterprise datacenter may be running multiple classes of applications, on hundreds or even thousands of systems. Administrators may schedule modifications almost daily to reconfigure systems, perform upgrades or apply patches. Depending on the processes that are in place, every change could potentially involve downtime.

Due to their ubiquity, Microsoft products are constantly under attack by hackers. In response, Microsoft frequently issues new security patches for its operating systems or for Microsoft Exchange, to close the door on a virus or to modify the rules for opening attachments. Although patches and upgrades are applied at the administrator's discretion as often as once a week, planned downtime could still regularly occur.

Performing these routine operations without disrupting service is a major concern for companies today. The immediacy of online services has produced a culture of end users who are fast approaching zero tolerance for delay. For online businesses, for example, an outage or slowdown that lasts as little as 5 minutes is almost guaranteed to drive impatient customers to competitive sites for goods and services. If they

experience the delay again, customers could be gone for good. From the standpoints of revenue and reputation, these are serious concerns.

LEVERAGING HIGH AVAILABILITY SOLUTIONS

Data protection and high availability solutions can be leveraged for planned downtime operations, dramatically shrinking both administrative costs and outage windows. These tools automate procedures to make administrators more efficient, reduce the possibility of human error and accelerate processes that could negatively impact end users.

The term *planned downtime* is in fact almost an oxymoron today. The downtime associated with most planned operations that could potentially affect uptime can be virtually eliminated. For example, volume management technologies can move data from one storage array together and allow all maintenance tasks to be performed online. Automated real-time replication of volumes can move data transparently and vendor-independently to other servers at any worldwide location. Automated failover with clustering technologies can move a live application to another server with minimal service interruption or move an entire datacenter to another location. Let's look at some typical examples.

Scenario One: Migrating an Application. Say you want to move all your users to another server. There is no usable server available within the datacenter, so you have to move the application, with all its user groups and all its data, to a server in a datacenter that is two time zones away. Volume management and replication software can mirror the data to a storage volume in the distant datacenter, which can dramatically reduce potential data loss and speed recovery time by making current data available instantaneously at an alternate location. Clustering software can move the application to the new server in a failover operation that preserves the state of the application and its user data. Clustering technology protects your applications both at the primary site and multiple wide area sites. These planned, automated operations can carry out the migration with virtually no impact to service lev-

els and allow users to continue to access their data and application without incurring the downtime normally associated with the migration of applications.

Scenario Two: Upgrading Microsoft Exchange. You're running Exchange in a non-clustered, non-automated environment, and you need to upgrade the system with a new service pack. You are forced to shut down the server and turn off all your users, cutting off access to arguably their most critical application. You then have to load in the new service pack, load in the new application, point the application to the data, and redirect your clients to the new server. If on the other hand, you're running Exchange in a clustered environment, performing upgrades has minimal impact on users. With automated clustering and volume management tools, copying the data to available disk and migrating the Exchange application service to another server within the cluster is a push-button operation saving hours of downtime.

Scenario Three: Server Consolidation. You have a number of standalone systems in your environment that are not generating the ROI they could. The decision is made to bring in a new class of server and storage hardware that can be configured with multiple domains. Using the clustering technology that you acquired as part of your availability solution, you simply add each standalone system to the new cluster and execute a migration command to move the application over. Your consolidation ratio is 12:1 (in this hypothetical case), and you wind up with a much smaller footprint for those applications, translating to higher reliability and availability. You can perform the same amount of work with fewer servers and reduce your hardware and maintenance costs.

Scenario Four: Performance-Based Migration. You have applications running in an environment where performance is degrading. The time required to process a transaction, to respond to a Web-page click, has reached the point where your monitoring software is sending alerts. Your disaster recovery solution includes a cluster server and cluster file system that enable you to migrate the applications, permanently or temporarily, to systems that are better able to handle the service. The process is automated. You eliminate downtime while maintaining application performance.

LOOKING AT THE BOTTOM LINE: SPEED, EFFICIENCY AND COST REDUCTION

The clustering, replication and volume management tools that make up an effective high availability/disaster recovery solution for many companies can contribute handsomely to ROI and provide real business value when they do double duty as facilitators of planned downtime projects.

One obvious advantage is that this is an excellent use of current assets since these technologies are already in place at most companies. No additional expenditure is necessary.

Another powerful benefit is that these tools automate procedures, reduce administrative costs, increase efficiency and eliminate the possibility of error. Many highly automated routine procedures can be initiated remotely by administrators or run as timed events at night.

The most significant benefit, however, is the ability of these tools to reduce downtime to a matter of seconds. Using migration technologies, users can be migrated to another set of systems without the need for them to reconnect, thereby significantly reducing the downtime associated with planned maintenance. The connection is persistent, and the state of the application is maintained, even for users who were conducting transactions.

Unforeseen events and site outages will happen. However, clustering, replication and volume management technologies can be leveraged to ensure the availability of data and applications, minimize the impact of failures on the business and ultimately align IT with business operations while significantly increasing the enterprise's ROI. 🐼

As Senior Manager, Mrs. Atwood is responsible for the high availability product line and disaster recovery solutions. Prior to joining VERITAS in 2001, she served as a product manager at Sun Microsystems. Mrs. Atwood holds a Bachelors Degree in Political Science from the University of California in San Diego and an MBA from San Jose State University. She has been published in VAR Business magazine, Dell Power Solutions Magazine, Nth Generation, Computerworld, Continuity Insights and Disaster Recovery Journal.

