Resonant
Communicatons

Virtualization:
Improving Disaster Recovery for SMBs

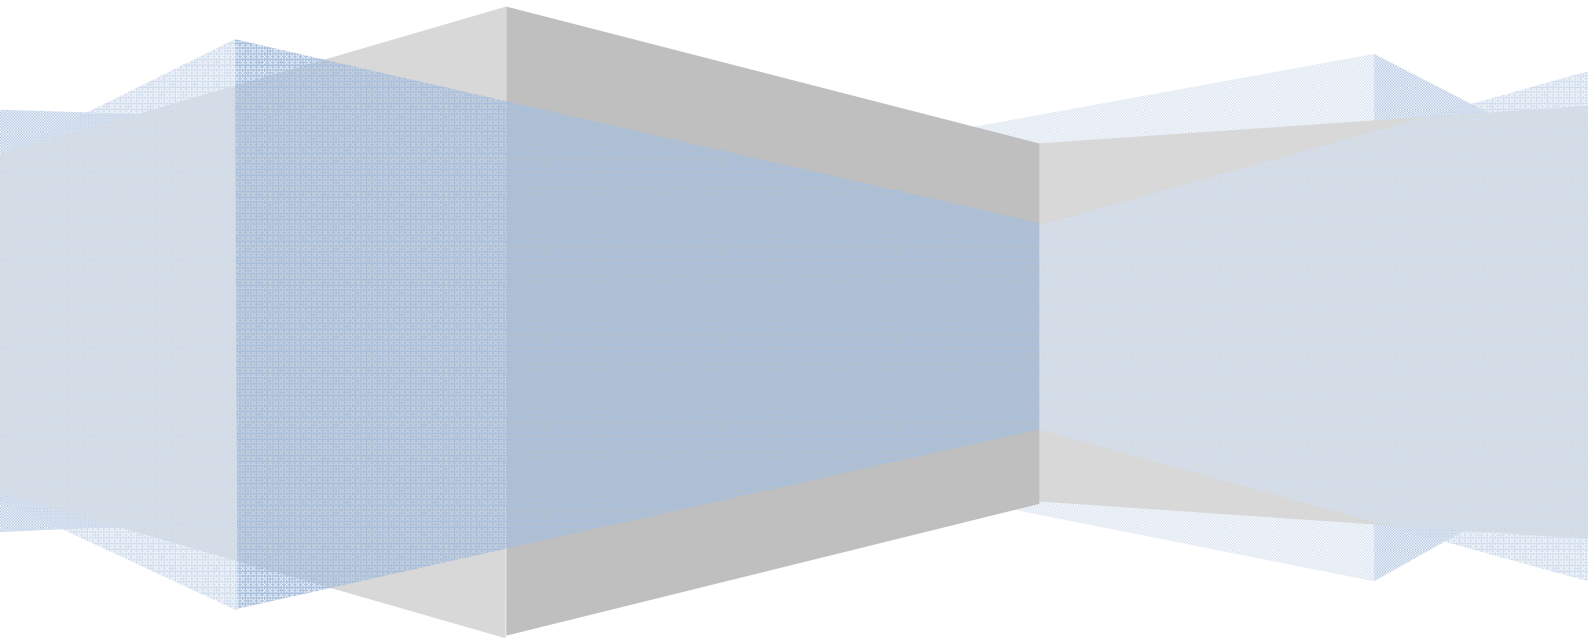

By Kerry Doyle

Improving Business Continuity & Disaster Recovery for SMBs

How Virtualization Can Improve Disaster Recovery

No matter the size of your business, in today's economic climate, having an IT infrastructure that includes cost-effective disaster recovery (DR) and business continuity (BC) solutions is a necessity. Virtualization now offers an affordable solution for small and midsize businesses (SMBs) to not only save on energy costs and space requirements, but to perform effective disaster recovery.

Traditionally, disaster recovery has relied on data backed up and stored on physical formats, such as tape, Blu-ray disc or DVD. Periodic backups are then archived for long-term storage or regulatory compliance. After a disaster, data can be restored to either the repaired hardware or to hardware located within an entirely new datacenter.

DR in a virtualized environment provides a range of options for recovering data and performing rollbacks to earlier, damage-free iterations of a primary virtual disk. On the other hand, BC comprises a more active, agile protection mechanism. It enables IT to recover files while simultaneously maintaining application processes with minimal downtime. Then, when disaster recovery is performed due to a malfunction or system failure, the process is generally transparent to users, who can continue to work without interruption.

Protecting business critical systems and data is crucial for providing SMBs with a competitive edge. The following article explains how virtualization can significantly improve both disaster recovery and business continuity.

Assessing the value of data

Since BC and DR implementations can be expensive, it's important to calculate the value of the data you wish to protect and the cost to the organization if this data was permanently lost. Otherwise, you risk overspending and committing resources to reclaiming data that's not worth the cost of recovery.

For example, if your company possessed 100,000 pieces of user data worth $10 per user, the value of this data would equal $1 million ($10 x 100,000). However, if the data was worth only $0.05 per user, the value would be $5000 ($0.05 x 100,000). It's clear that committing both physical and financial resources to recovering data worth $1million would be crucial to the long-term viability of this SMB.

Data Retrieval

Generally, once a decision is made to retrieve data, the expense for an SMB is based on the tools implemented. The type of disaster, from small application failures to large data center malfunctions (see sidebar: Disaster Types and Virtualization Solutions), often determines the resources that need to be added to an SMB's existing set of software applications and hardware.

They can include:

- Hardware for a cluster (i.e., computing resource, network infrastructure and storage)
- Cluster licenses
- Additional software licenses per node
- Operating System licenses
- New security solutions across applications and data centers

In addition, an SMB should calculate the number of IT resource hours and consulting expenses required for installation, configuration, maintenance and testing of the new solutions, adding these to the overall expense of disaster recovery.

Finally, part of any DR test should be the adequate restoration of data. This process in itself can be quite time consuming. For example, with a standard Windows Server, a number of steps are needed, from re-installing Windows and final reconfiguration to thoroughly testing applications and returning them to the production server.

The problems related to quick data recovery after a disaster can be compounded by the added pressure on an IT administrator to perform tasks correctly and efficiently while trying to function within a disaster situation.

Simplifying recovery

Virtualization can help alleviate much of the pressure, time and added costs of DR and BC. One way virtualization does this is by providing built-in clustering tools, such as high availability. An optimal high availability solution is where the data and operation of a cluster of physical servers at a primary site are replicated in an identically configured cluster of stand-by virtual machines at another backup site.

Virtualization also provides virtual machine replication across a Wide Area Network (WAN). Thus, when the primary site malfunctions, data center operations can be quickly transferred via the WAN to a stand-by site in a process called *failover*.

SMBs will notice that when you virtualize a physical machine, there is an immediate improvement in

availability. Improvement is realized because hardware and firmware revisions are not required. In other words, the need to apply patches or re-embed typically small code sets found on hardware, such as motherboard BIOS updates, has been alleviated.

Virtualization-based backups

In addition to the protection of storage subsystems, the ability to seamlessly rollback changes to a point before a disaster occurred as well as to create dependable backups are two of the most important aspects that virtualization offers.

- With virtualization, disks are abstracted so only a few files are needed for comprehensive backups. A virtual machine can then be recreated using one or two large virtual disk files and meta-data files.

- Snapshots of disks made before the installation of patches are useful for displaying where changed blocks have been written. Administrators can then rollback changes to the primary virtual disk, still in its pristine and unchanged condition.

- Virtualization consists of running multiple OSes on the same server & offers scripts at installation for each guest OS.

These scripts have the means to ensure that data on the virtual disk is quiescent, or at rest, which leads to better backups. These tools also allow for file level restore, in case the data file becomes inadvertently damaged.

Versatility with Fault Tolerance

One disaster recovery tool available to SMBs mentioned earlier is high availability (HA). Based on the linkage of two or more servers, any failure in this configuration allows the workload to be immediately transferred to another operating server. While unplanned downtime is minimal, users and applications do notice these interruptions and work-in-progress is lost. Fault Tolerance, on the other hand, is the ultimate tool for both business continuity and disaster recovery for SMBs. The tool resides within a cluster of virtualization hosts and provides a method for almost instantaneous failover to a functioning “shadow” server from a failed server.

In this scenario, downtime is almost completely eliminated as well as adverse effects to users and applications, and entails no data loss.

There are several additional tools that help to improve DR. These virtualization-savvy backup tools work with the virtualization host to safely extract a backup as well as to replicate data from one site to another. These can be helpful in creating prepared host servers, or hot sites. They also allow for the relocation of virtual machines between geographically disparate datacenters.

Conclusion

Virtualization offers an effective solution for DR and BC for SMBs over traditional methods. For cost effectiveness, data valuation is critical. SMBs must also take into account additional hardware, software and IT considerations. However, virtualization can decrease expenses enabling companies to sustain critical operations while meeting 24/7 demands. Regardless of the disaster, from application failures to data center malfunctions, virtualization enables SMBs to maintain their competitive edge and to avert potential crippling disaster scenarios.

Improving Business Continuity & Disaster Recovery for SMBs

Disaster Types and Virtualization Solutions