

Redmond

THE INDEPENDENT VOICE OF THE MICROSOFT IT COMMUNITY

REDMOND IN-DEPTH REPORT:

Disaster Recovery and the Question of Balance

Because one size does not
fit all, smart enterprises
consider a litany of factors.

BY DAN KUSNETZKY



DR is really a combination of processes—risk analysis, planning, implementation and ongoing operation—combined with both hardware and software.

Disaster Recovery (DR) is a hot topic today. Despite that fact, one of the main challenges to a successful DR implementation is knowing exactly what it is. Vendors use DR as a blanket term when speaking about many different types of products, even though each of these products is doing something different.

One reason for confusion is that DR is really a combination of processes—risk analysis, planning, implementation and ongoing operation—combined with both hardware and software. They're jointly designed to respond to some sort of disaster quickly, reliably and efficiently, allowing business operations to continue.

What Constitutes a Disaster?

Disasters include a wide range of events, including:

- Natural disasters such as a fire, flood or storm.
- Hardware failures like power, air conditioning, systems, system components, storage networks, storage devices, network connections and network devices.
- Software failures such as poorly implemented applications, database failures, loss of messaging from one software component to another, or application-framework failures.
- Security issues like malicious injection of SQL code, corruption or loss of files.
- A wide range of human errors that can bring down even the best-designed application environments.

It's easy to see that planning for disasters involves many levels of business, facilities and IT management, as well as experts in systems software, virtualization technology, application frameworks, application development, database management, storage and networking.

I'm going to focus on the IT hardware, software and services elements. A quick examination of those areas reveals that vendors offering products that touch on any aspect of "keeping the lights on" will claim the full mantle of DR, even though the use of its product or

service isn't a complete answer. These vendors often use catch-phrases such as "continuous processing," "always on" or "nonstop."

The Right Tool for the Right Job

Consider many of the different ways vendors address DR. Although few actually address all of an enterprise's requirements, each of the following approaches are presented as if they're really the whole enchilada, rather than just a few tortilla chips:

- Hardware components that support continuous processing; that means nonstop, fault-tolerant computers, such as the ftServer from Stratus Technologies. This approach focuses on the underlying host systems and assures that end users will never see a failure. These systems are designed with multiple layers of redundant hardware and special firmware that detect failures and move processing to surviving system components. Failover takes only a number of microseconds and is automatic.
- Clusters of systems designed to detect slowdowns or failures and move applications and data to maintain continuing operations. Suppliers such as Dell Inc., Hewlett Packard Enterprise (HPE), IBM Corp., Microsoft, Oracle Corp., Red Hat Inc., SUSE and many others offer this type of DR solution. Cluster-software managers monitor the health of systems, applications, and application components, and move functions to another system when a failure or slowdown is detected. Applications typically must be designed to work with this cluster-software manager. Failover can take hours, depending on the design of the cluster manager. Systems typically are supporting workloads, and are there as warm standbys for all other systems.
- Systems that house backup software. These are appliance servers pre-loaded with backup software. In this setup, applications and data are constantly backed up to either storage in the datacenter or to a cloud-storage service. Upon failure, this data can be manually or automatically recovered. While some products can detect failures and start a recovery process, many require manual intervention. Full recovery can require a number of days, depending on the complexity of the environment.

While some products can detect failures and start a recovery process, many require manual intervention.

If a slowdown or failure is detected, applications and their components are moved to a surviving system.

- Storage systems that keep multiple copies of each data item, making it possible for applications to continue accessing and updating these data items even though a component has failed. Suppliers such as EMC Corp., Hitachi Data Systems (HDS), NetApp and many others include this capability in their storage servers. Replication software that can keep copies of data items in several places is available. Replication occurs either in the storage server itself or in host systems attached to the storage server. In the case of a failure, operations staff can point applications to data items in another location. Some products will automatically redirect storage requests, rather than requiring manual intervention.
- Storage software that, like storage hardware, keeps multiple copies of each data item. This approach, offered by suppliers such as DataCore Software, Citrix Systems Inc./Sanbolic and others, also make it possible for the data to be replicated to other datacenters or cloud services.
- Software that supports continuous processing. This can include special-purpose virtual machine (VM) software such as everRUN from Stratus Technologies. This approach is very similar to a combination of continuous-processing hardware combined with monitoring, management and application migration. If a slowdown or failure is detected, applications and their components are moved to a surviving system. The detection and migration process is automatic and simulates how continuous-processing systems work. This often requires multiple systems to be used as hot or warm standbys.
- Software that monitors VM operations and initiates a migration from one host to another system. This category includes VM monitoring, management and migration tools offered by Citrix, Microsoft, Red Hat, SUSE and VMware Inc. When the monitoring software detects a slowdown or failure, applications or even entire workloads can be migrated from host to host, datacenter to datacenter, or even from datacenter to the cloud.

‘One-Size-Fits-All’ Is a Myth

It’s clear from reviewing these different approaches that each has its benefits and its limitations. Some approaches provide an environment

that never fails, but at a very high cost. Others are less costly, but the failover process can take some time, possibly resulting in data loss.

Three key questions must be considered for each application and application component for the proper mix of technology to be selected:

- How much money does the enterprise want to invest in availability for specific applications and their data?
- How quickly must each application become fully available and functioning?
- How much overhead can be supported to provide availability?

Depending on the answers to those questions, a mix of solutions can be selected. Here are a few suggestions:

- Hardware-based solutions, such as continuous-processing systems combined with either storage servers or storage software, can be deployed for the most-critical applications. The type of storage, which includes rotating media, flash-based devices or cloud services, can then be selected based on the cost, performance and availability parameters that are best for the application.
- Software-based solutions, such as the use of VM migration combined with storage-software solutions, often cost less and are more flexible. Remember, however, that they're often slower than hardware-based solutions. Less-critical applications can be hosted on these types of solutions.
- Cloud-based solutions, such as those offered by nearly all cloud services providers, often appear to be the least complex and lowest cost. The failover process, however, can be lengthy, and access to data items can be slow when compared to local storage.

Software-based solutions, such as the use of VM migration combined with storage-software solutions, often cost less and are more flexible.

Another important point to consider is storage performance and its impact on each of these potential solutions. Here are some general guidelines:

- In-memory storage is used for some extreme transactions, Big Data and Internet of Things (IoT) applications. This approach

In the end, there is no single answer that addresses the needs of all enterprise applications.

offers the highest performance for the highest cost. They also typically require backup software to be used to prevent data loss in the case of system failure.

- Flash-based storage solutions offer high levels of bandwidth and low levels of latency. While very fast, these solutions are costly compared to traditional media. Suppliers of all-flash storage systems, such as EMC, Kaminario, NetApp, HDS and others would point out that the cost per gigabyte has been dropping dramatically, and is now more comparable to other types of storage.
- Rotating media-based technology offers solutions to any but the most strenuous application and business requirements. The intelligent caching capabilities of storage-software products, such as those offered by DataCore or Citrix/Sanbolic, often provide similar levels of performance using a small amount of system memory or flash storage. When combined with other types of storage, they can offer lower overall cost.

How Long Is a Piece of String?

In the end, there is no single answer that addresses the needs of all enterprise applications. Each form of DR and storage technology fits the requirements of some applications and not others. This is a bit like being asked the question, “How long is a piece of string?” The only correct answer is to measure the piece of string in question.

Enterprises must take the time to understand their own application portfolio, along with their own business and availability requirements. This must be done for each application. Only then can the proper mix of approaches be selected to address both the need for availability and disaster tolerance, balanced against budgetary limitations. **R**

Daniel Kusnetzky, a reformed software engineer and product manager, founded Kusnetzky Group LLC in 2006. He's literally written the book on virtualization and often comments on cloud computing, mobility and systems software. He has been a business unit manager at a hardware company and head of corporate marketing and strategy at a software company. In his spare time, he's also the managing partner of Lux Sonus LLC, an investment firm.