CONTENT PROVIDED BY

Redmond
THE INDEPENDENT VOICE OF THE MICROSOFT IT COMMUNITY

SPONSORED BY

DELL Software

Disaster Recovery and High Availability

These two solution categories sometimes seem at odds with each other, but they're both alive and well. Find out how in the stories that follow.

Disaster Recovery: Not Dead Yet

In the age of virtualization, many think high availability trumps disaster recovery. Think again: It's more necessary than ever.

BY JON WILLIAM TOIGO

One doesn't need to look very far to find hypervisor vendor marketing materials, analyst opinions and even a few trade press articles declaring the death of disaster recovery planning. Most are following a line of reasoning from the server virtualization evangelists to make the case that building server hosts with high availability (HA) in mind—using active-passive failover clustering models, for example—obviates the need for disaster recovery planning.

Most disasters are local in their cause and effect.

Add to the fact that most disasters are local in their cause and effect (surveys suggest that 95 percent of IT outages result from application and software glitches, hardware component failures, malware and viruses, and scheduled maintenance), and a business-savvy question must be asked: Why spend a lot of money to build the capability to recover from a bigger disaster—a severe weather event, large-scale infrastructure outage and so on—that has a low statistical probability of

happening, anyway? It's not hard to think of other things to do with budget dollars than invest them in an insurance policy that may never need to be used.

Failing To Plan...

In reality, those arguments carry on a tradition of anti-planning sentiments dating back to Noah and the Ark. They tend to hold sway until one of two things happens: an "every-200-year" event like Hurricane Sandy (only to be followed by a second once-in-a-lifetime event the following year), or when a new regulatory or legal mandate requiring business continuity planning and data protection and preservation comes into force.

The problem is that few organizations have any personnel on staff who are acquainted with even the fundamentals of disaster recovery (DR) or business continuity planning (BCP).

When either of those things occurs, senior management tends to take a greater interest in the preparedness of the firm for a disaster, and tends to allocate budget to the planning project.

The problem, however, is that few organizations have any personnel on staff who are acquainted with even the fundamentals of disaster recovery (DR) or business continuity planning (BCP). Some have been content to accept their vendor's word for it that deploying two virtual servers in a failover cluster constitutes "business continuity," though ISO Standard 22301, covering business continuity, defines the term quite differently.

Is the nomenclature important? Very much so, especially if your organization is claiming to adhere to the ISO standard as a demonstration of compliance with a regulatory requirement, such as the Health Insurance Portability and Accountability Act (HIPAA).

If patient health-care data is lost owing to a disaster that could have been prevented with basic data protections that would have been part of an ISO 22301 program, the health-care company could be on the hook for regulatory non-compliance. At a minimum, that might earn some unflattering coverage on the front page of The Wall Street Journal; at worst, it might open the door to countless lawsuits, not only for compliance issues but also for fraud.

Anonymous 1s and 0s

A real business continuity program focuses, as the name implies, on the business—or, rather, on business processes. The planner

undertakes a process to assess the criticality of each business process, which includes developing an idea of what processes require fast restores, and which can wait a while in the wake of an interruption event.

With process criticality determined, the planner locates the applications that support the business process and discovers the data associated with those applications. Data inherits its criticality from the process it serves. Without business context, data is just a bunch of anonymous 1s and 0s.

The analytical process concludes with a discovery of where and how data is hosted. It's necessary to find where the data physically resides to help determine the most efficacious way to apply protective services to the data. And, of course, planners also need to understand rates of change in data (how often it's updated), and rates of data growth to identify appropriate backup and recovery strategies.

It's necessary to find where the data physically resides to help determine the most efficacious way to apply protective services to the data.

A business process focus (see **Figure 1**) is essential for business continuity planning. It's the process, after all, that needs to be recovered following an unplanned interruption event. So, as onerous as it sounds, planners need to develop a clear understanding of the

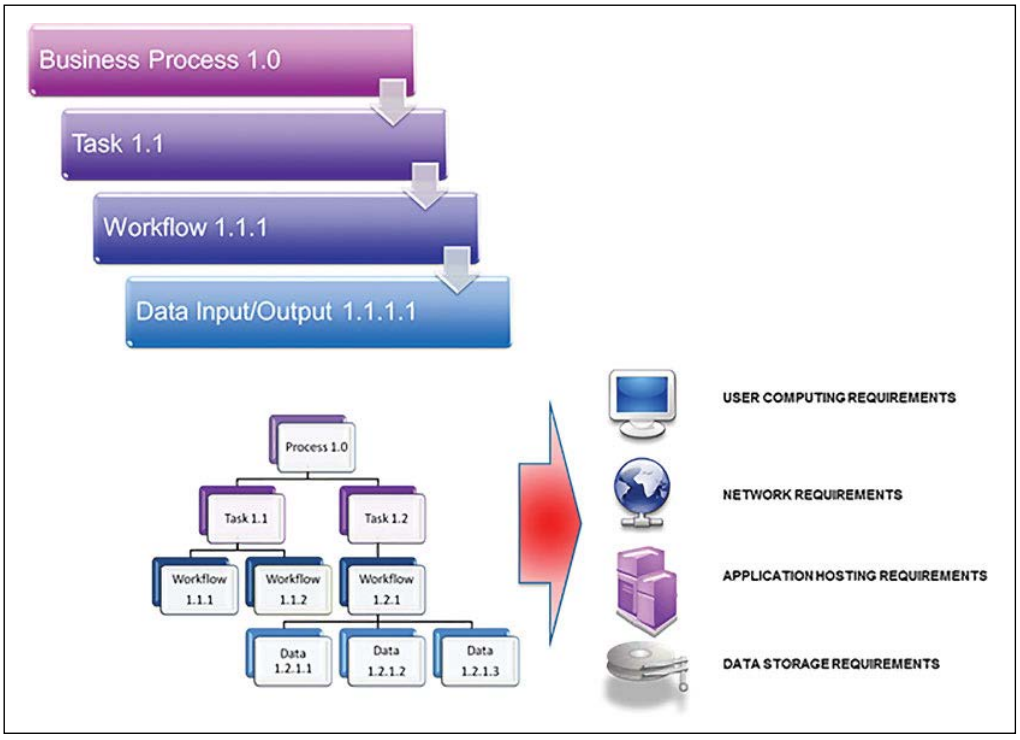

Figure 1. *A typical business process workflow.*

criticality of each business process. This investigation also identifies all relationships between business processes and interdependencies between their applications and data (which often are not intuitively obvious); a key requirement for building an effective recovery strategy.

Business process criticality drives recovery time objectives, which, together with budgetary constraints, define recovery strategy options. The criticality assessment may also have other value: Because it's likely the first time since systems were first rolled out that anyone has sought to document the alignment of the business with its IT infrastructure, the information collected can be of enormous value not only in DR/BCP, but also compliance planning, archive planning and security planning.

Business process criticality drives recovery time objectives.

Stretching Your Cluster over Distance: Don't Forget About Einstein

Dream as we might about WAN-based "stretch clustering"–an extension of high availability (HA) architecture beyond the facility walls to another office or to a cloud services provider–we must consider the practical constraints imposed by latency (distance-induced latency occurs because data cannot be moved across an interconnect at faster-than-light speed) and jitter.

Distance-induced latency is encountered as the length of an interconnect–a metropolitan area network (MAN) or wide area network (WAN)–exceeds 70 km (see **Figure A**). Latency translates

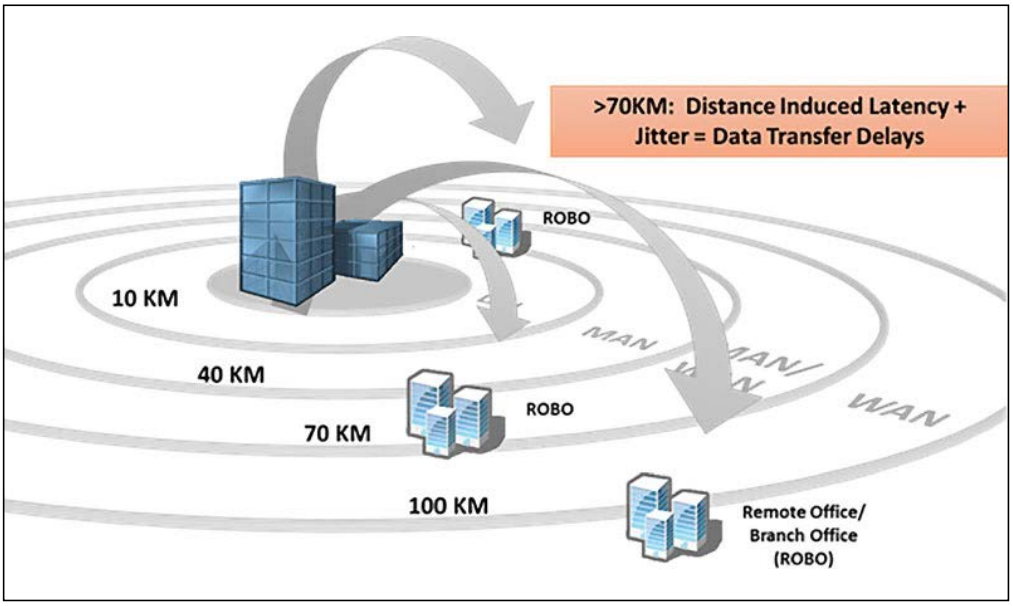

Figure A. *Distance-induced latency and "jitter" affects disaster recovery.*

into a delta or difference in the state of data at the source storage and at the target that may not be important when transporting a copy of backup data to a branch office or a cloud service, but that can be very problematic when trying to synchronize geographically dispersed servers running interdependent transaction-oriented applications.

In addition to latency on public networks, planners also need to consider the impact of jitter. **Figure B** identifies key types of jitter and their causes. Jitter also impacts the efficiency of data replication and remote application operation.

Some key points are often lost in translation during discussions of stretch clustering and data replication over distance:

Data de-duplication and compression doesn't make data move more quickly through a network interconnect.

- Data de-duplication and compression doesn't make data move more quickly through a network interconnect. Think about a traffic jam: the SMART car moves no more quickly than the 18-wheeler.
- The 70km "latency wall" is fixed: No amount of link optimization will change Einstein's speed of light rules. Also, 70km isn't always "as the crow flies": in metro areas, a lot of distance is consumed by circuitous wiring paths, which must go around obstacles, up elevator risers and so on.

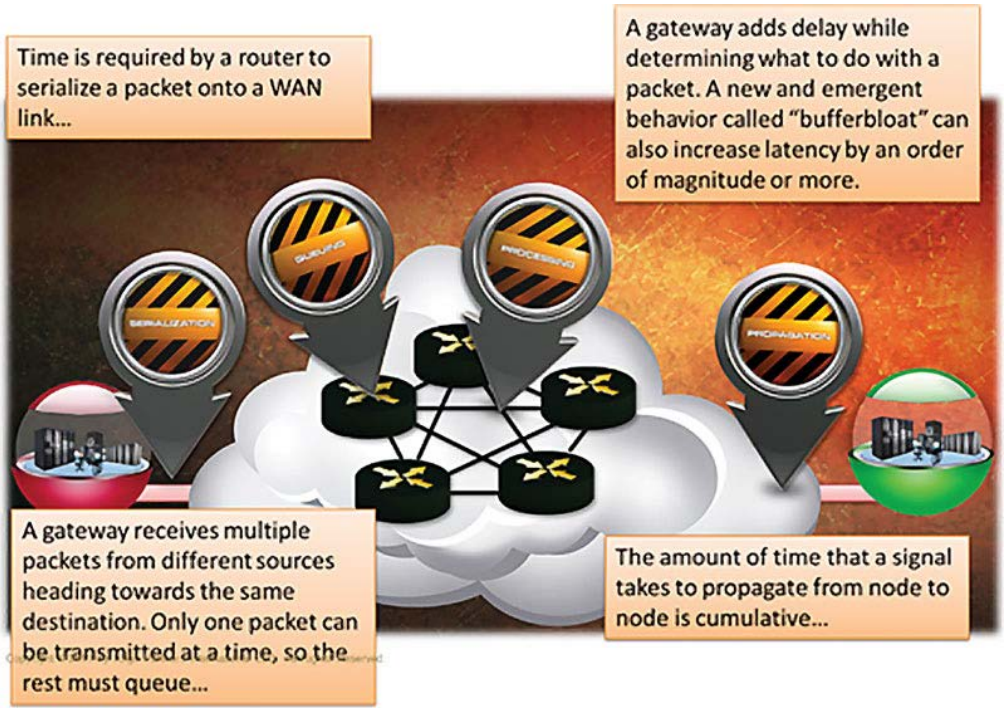

Figure B. *Causes and impacts of "jitter."*

- You need to know exactly where a cloud DR service is located. The great connection speed you're seeing might just mean that the physical facility is right across the street, affording no meaningful DR protection at all.

Impact Analysis

This "impact analysis," as some call it, is the heavy lifting of DR/BCP. It's the only way to do an effective job of building strategies to protect data assets in a granular way that will support the organization's recovery priorities. It provides the only basis for fitting the right recovery strategy to the right recovery target, based on recovery time objectives.

There has never been one all-encompassing strategy for recovery.

There has never been one all-encompassing strategy for recovery; planners have always selected a spectrum of options for data and application re-hosting. Since the earliest days of mainframe computing, the options for safeguarding business application processing ranged from laissez-faire approaches (just take a backup and, if a calamity occurs, work with your vendor to resupply whatever hardware that has been lost), to full redundancy (what we call active-active clustering over distance with automatic failover to the remote site if the primary becomes compromised). **Figure 2** shows the range of options.

Clearly, redundancy (HA) provided the shortest recovery time, but it was also the costliest strategy to implement and maintain. This has limited its use.

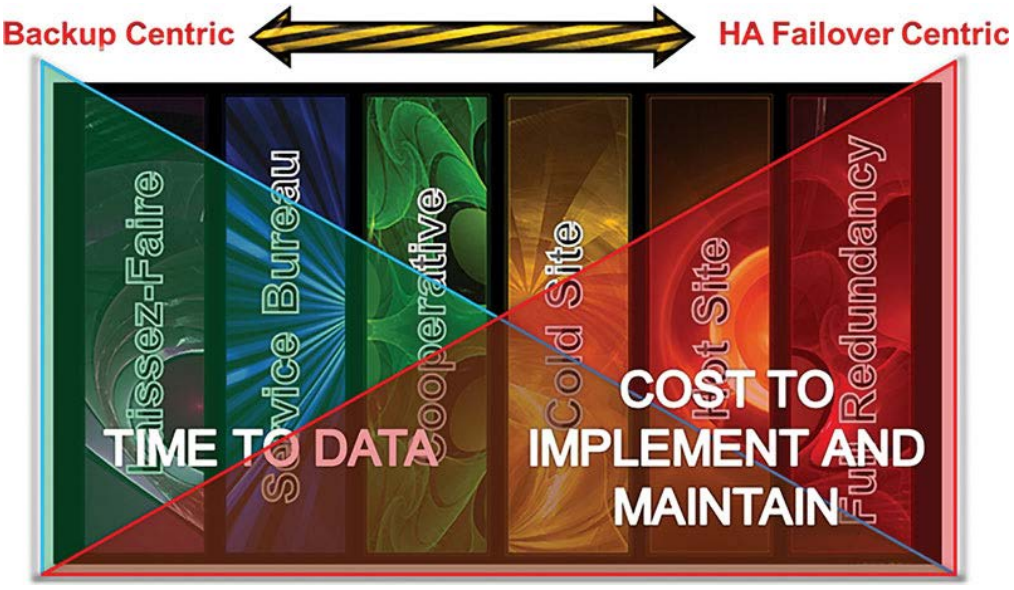

Figure 2. *Many factors go into a proper backup/recovery and business continuity plan.*

Failover Clustering

As mentioned previously, there's no "one size fits all" strategy for DR or data protection, at least not one that fits all protection requirements well or cost-effectively. That goes double for failover clustering in the virtual server world.

In truth, organizations don't need, and can't really cost-justify, failover clustering for every application. Following most outage events, companies report that their focus was on recovering about 10 percent of their most mission-critical apps. The other 90 percent waited hours, days or even weeks to be recovered, without devastating consequences to the firm.

There's no "one size fits all" strategy for DR or data protection.

Clustering 'Super-Processes'

Noting that HA clustering is not a cost-effective overall strategy, however, doesn't necessarily resonate with IT folks who have been inundated with hypervisor vendor claims about their "panacea" technology for built-in business continuity. Indeed, many virtualization administrators don't clearly understand the challenges of clustering or the prerequisites for making a failover successful.

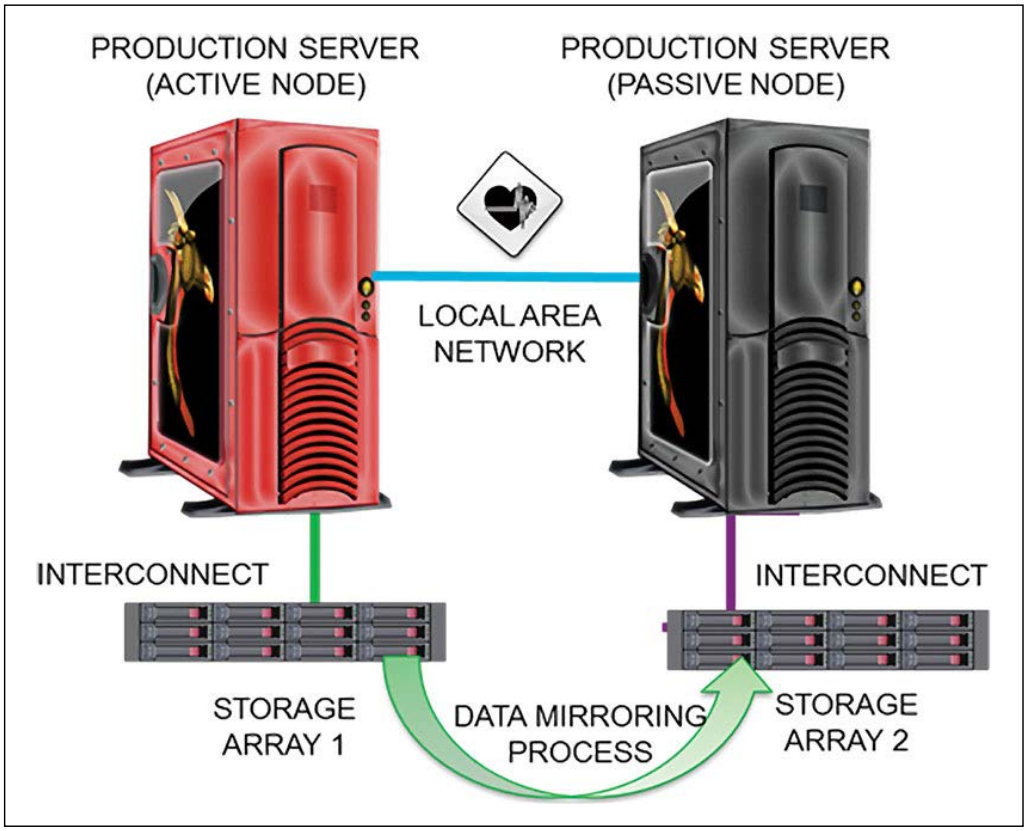

Figure 3. *A typical heartbeat-based, active/passive failover cluster.*

Clustering requires some pretty heady provisions to be successful. It actually requires the careful implementation of two "super-processes." One is an ongoing software-based communication across a physical interconnect (usually a LAN) between two or more server nodes that essentially serves as a "heartbeat" monitor, continuously checking the health of the communicating nodes. If this heartbeat fails for any reason, the failover of workload from the active to the passive node is typically automatic (see **Figure 3**, page 7). (This is slightly different in the case of active-active clusters, in which workload is balanced between nodes until a node fails and the entire load shifts to the cluster partner.)

The second super-process is ongoing data replication between the storage supporting each clustered server. A mirroring process is typically set up between the storage connected to each server, so that the same data will exist behind each server should a failover process need to be initiated.

A big problem with the mirroring superprocess is that mirrors are seldom–if ever–checked.

Cracked Mirrors

A big problem with the mirroring superprocess is that mirrors are seldom–if ever–checked. Doing so requires that business apps generating data be quiesced, that caches be flushed to the primary write target storage, that the data then be replicated to the mirror storage, and that the mirroring process be shut down, just so you can look at both the primary and mirror volumes to check that both primary and replica data are the same.

Once verified, the mirror and business apps need to be restarted and re-synchronized, which can be a dicey process with career-limiting consequences if re-sync fails. Bottom line: mirrors mostly go untested, so there's no certainty that the data required to "seamlessly failover" from one host to the other is actually present on both systems.

Even with these technical and procedural challenges, failover clustering with mirroring is the prescription of most hypervisor vendors for ensuring continuity of workload processing and data availability. Indeed, this model is recommended by some vendors to ensure business continuity in the face of facility-wide disasters, in which recovery at an alternate location may be required. This "stretch clustering" or "geo clustering" strategy is also fraught with challenges and must be considered carefully.

Further complicating the world of virtual server failover is the ability of VMs to migrate from one physical host to another for purposes of load balancing, optimizing infrastructure resources or freeing up gear for maintenance. While potentially beneficial to datacenter agility, this capability might complicate the clustering/mirroring processes even more, turning it into a veritable shell game and growing the number of data replicas to the point where storage capacity demand accelerates beyond all acceptable rates.

Skyrocketing Capacity Requirements

The impact of disaster recovery strategy on storage capacity demand growth is a big concern in many firms. At confabs last summer, IDC analysts were revising their 40 percent per year capacity demand growth estimates upward to 300 percent per year in highly virtualized environments. This partly reflects shifting workloads, but also the "minimum three storage nodes behind each virtual server" configurations now being promulgated by VMware Inc., Microsoft and others as part of the trend toward software-defined storage models. The next month, Gartner Inc. doubled the IDC estimate, noting that it didn't take into account DR backups or other data copies.

The impact of disaster recovery strategy on storage capacity demand growth is a big concern in many firms.

Clearly, whatever strategies are selected for data protection and DR/BCP also need to respect available budget and collateral costs. Strategies also need to reflect the real world.

Given all this, it's clear that claims of traditional DR planning being relegated to the dustbin of history by HA architectures embedded in virtual server computing are wrongheaded and foolish. If anything, the advent of hypervisor computing has brought about an increasingly siloed IT infrastructure with proprietary hardware and software stacks organized under different vendor hypervisor products. This is having the effect of creating multiple, separate DR targets, each in need of its own strategies for protection and recovery.

Storm Warning

In the contemporary enterprise, where some apps continue to operate without hypervisors and where the virtualization environment features multiple hypervisors with separate server, network and storage components, the need for a robust DR/BCP program has never been greater. Truth be told, HA has always been part of the spectrum of strategies available to planners; but it introduced

As we enter the new severe storm season, and the potential for data disruption it represents, it's probably wise to remember these points.

complexities and costs that found it's used only behind the most mission-critical applications, with the most demanding "always-on" operational requirements.

As we enter the new severe storm season, and the potential for data disruption it represents, it's probably wise to remember these points. R

Jon Toigo is a 30-year veteran of IT, and the Managing Partner of Toigo Partners International, an IT industry watchdog and consumer advocacy. He is also the chairman of the Data Management Institute, which focuses on the development of data management as a professional discipline. Toigo has written 15 books on business and IT and published more than 3,000 articles in the technology trade press. He is currently working on several book projects, including The Infrastruggle (for which this blog is named) which he is developing as a blook.

High Availability: Past, Present and Future

To understand what HA solutions best fit your environment, you need to understand their history and how they've evolved. BY DAN KUSNETZKY

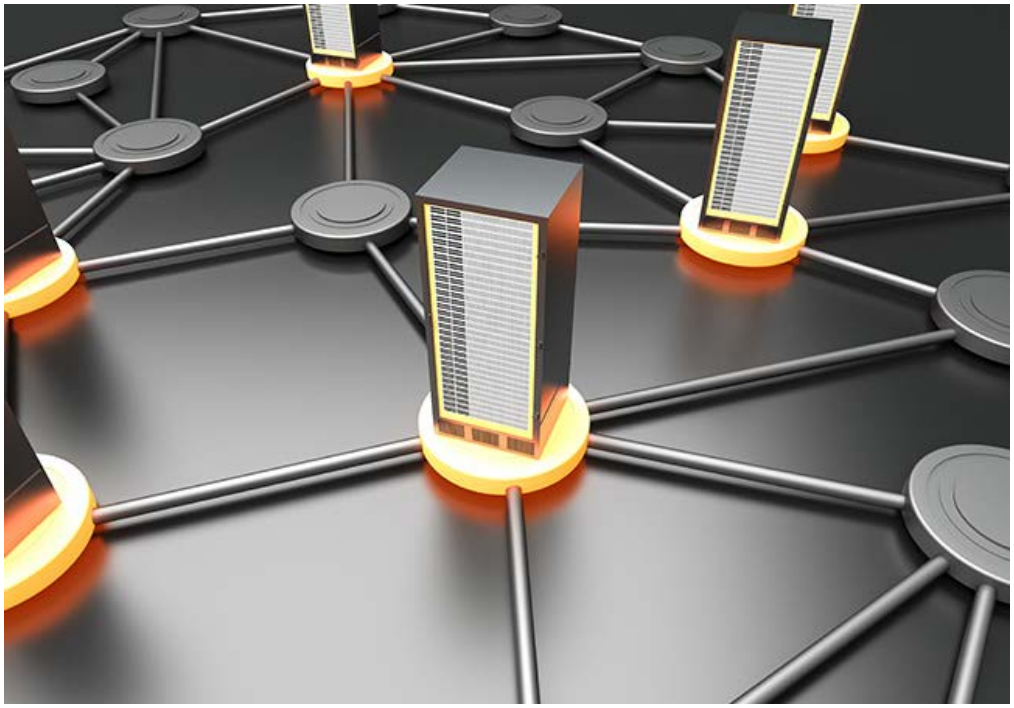

Most HA solutions rely on redundant hardware and special-purpose software designed to make the best use of that hardware.

High Availability (HA) is a topic with a great deal of history. Different approaches have been used over time to make sure applications, services, databases, networks, and storage remain available and reliable to support enterprises. As enterprises have grown increasingly reliant on information technology-based solutions, the need for these solutions to always be available has increased as well.

Most HA solutions rely on redundant hardware and special-purpose software designed to make the best use of that hardware. Virtualization and cloud computing are upending earlier approaches to HA. Organizations have learned that the use of virtualized access, applications, processing, network and storage makes the creation of HA solutions easier. They've also learned that virtualization makes it easier to use off-site cloud hosting as part of an HA solution.

HA solutions can be expensive, though, and an enterprise's portfolio of IT solutions might not need the same level of availability. Business-critical functions are likely to need the highest levels of availability, while the requirements for business support functions are not likely to be as high.

Enterprises would be wise to understand all of the following approaches to HA, and make the proper choice for each of their workloads.

Enterprises would be wise to understand all of the approaches to HA.

A Brief History of HA

When applications were more monolithic back in the 1960s through the 1990s, the UI, application logic, storage management, data management and networking functions were all hosted together on a single system. Back then, the industry focus was on making the systems themselves "fault tolerant."

This was accomplished by designing mainframe systems that used multiple processors, stacks of memory, storage adapters and network adapters; they included system firmware that monitored the health of individual components and moved workloads to surviving components in case a component failed or became unresponsive. IBM Corp. used "Parallel Sysplex," a special marketing catchphrase to describe these systems.

Parallel Sysplex failover took only a few microseconds or milliseconds. People using these workloads were usually unaware that a failure took place at all. These systems were extremely expensive when compared to standard off-the-shelf configurations, and were only used to host the most critical workloads.

IBM continues to make continuous processing mainframe configurations available today.

Suppliers such as DEC (now part of Hewlett-Packard Co.), Stratus Technologies (now owned by Siris Capital Group) and Tandem Computers developed similar technology in a smaller form factor–the minicomputer. IBM resold Stratus computers using the System/88 name.

As with the mainframe continuous processing systems, these systems were composed of redundant components and special-purpose firmware that detected failures and rapidly moved workloads so they could continue processing.

Failovers typically would only require milliseconds, and the users of these workloads were left unaware that a failure happened.

Because these systems were also quite expensive when compared to the off-the-shelf minicomputer competitors, they were only adopted to support the most critical workloads.

Failovers typically would only require milliseconds, and the users of these workloads were left unaware that a failure happened.

HP Integrity and Stratus ftServer systems are available today to address these business requirements.

Clustering

Suppliers hoping to address requirements for performance, reliability and availability worked to create more software-oriented solutions.

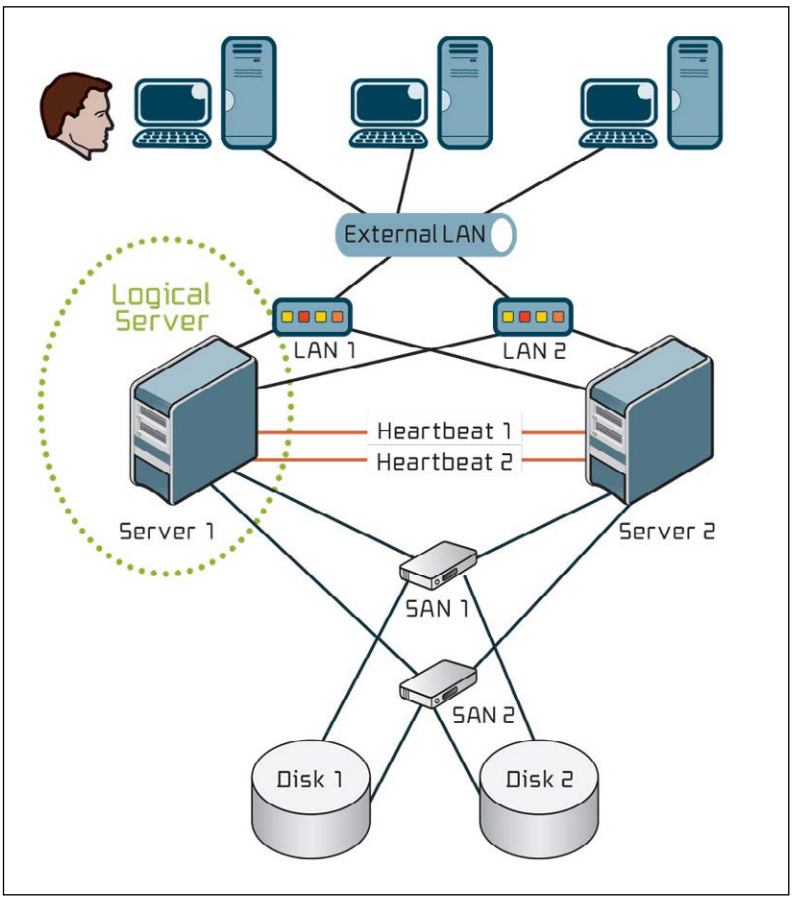

Figure 1. *A typical two-node High-Availability cluster.*

Rather than focusing on special-purpose hardware and firmware, these companies focused on special clustering and workload management software. The software orchestrated the use of either off-the-shelf networking solutions or special-purpose clustering networks.

Although clustered systems are likely to have been created by researchers as early as the 1960s, the first commercial offerings were the Datapoint ARCnet in 1977, which wasn't a commercial success, and the DEC VAXcluster in 1984, which was an overwhelming success and is still in use in many enterprises today.

Different layers of virtualization technology are deployed, depending on the goals of the enterprise.

These hardware configurations were used in a number of different ways. Each had a different goal and could be considered the earliest use of access, application, processing, networking and storage virtualization.

Customers deploy clusters, like those in **Figure 1**, page 13, to address the requirements for raw processing power, access availability, application availability, database availability, processing availability and even storage availability.

Different layers of virtualization technology are deployed, depending on the goals of the enterprise. Kusnetzky Group LLC has divided this virtual cake into seven layers, which you can read about at **VirtualizationReview.com/7LayerModel**.

Access Clusters

In access clusters, the basic cluster hardware configuration is used to make entire application systems available by using what is now thought of as "access virtualization" technology. Applications are installed on several cluster nodes, and if the node supporting the work of one group of users begins to fail, workload access is shifted from the failing system to one of the surviving cluster nodes.

While this appears similar to an application cluster, the failover and workload management is being done at the access level rather than the application level. Applications aren't aware of this technology and don't need special APIs or to be specially architected for this failover to occur.

This type of cluster relies upon data being housed on a separate part of the cluster devoted to storage access, on the storage services of another cluster, or on a storage-area network (SAN) so that data remains available even if the systems hosting the applications themselves failed.

Because access virtualization is the main virtualization technology in this type of cluster, application and storage hosts might be housed in the same or different datacenters.

Suppliers such as Citrix Systems Inc., Microsoft and VMware Inc. supply this type of technology.

Application Clusters

In application clusters, the basic cluster hardware configuration is used to make applications or application components available by using what is now thought of as "application virtualization" technology.

In application clusters, the basic cluster hardware configuration is used to make applications or application components available by using what is now thought of as "application virtualization" technology.

Application virtualization technology is used to encapsulate applications or their underlying components. The application virtualization technology controls access to these virtualized components. As users request the use of these applications, the workload management portion of this technology reviews the available processing capacity of the systems it's monitoring, selects a system to execute the application based on policies and the availability of processing capacity, and then starts up the application or sends the user's requests to an already-running application instance.

If an underlying system is failing, the user's workloads are automatically moved to another system in the cluster, or connected to workloads already running on another system.

While this appears similar to an access cluster, the failover and workload management is done at the application-component level. Applications must be architected to work with the application virtualization's workload management tool to enable workload monitoring, management and migration. So, unlike access clusters, the applications are extremely aware of this technology and must use special APIs or be specially architected for failover to occur.

This type of cluster relies on data being housed on a separate part of the cluster devoted to storage access, or on a SAN so that it remains available even if the systems hosting the applications themselves fail.

Access virtualization technology is often utilized, as well, so user access can be easily and automatically migrated from the failing systems to the new systems.

Application virtualization is the main virtualization technology in this type of cluster; storage hosts could be housed in the same or different datacenters.

Access virtualization technology is often utilized so user access can be easily and automatically migrated from the failing systems to the new systems.

Suppliers such as AppZero, Citrix, Microsoft, Novell Inc. and VMware offer application virtualization products today.

Processing Clusters

In processing clusters, the basic cluster hardware configuration is used to make entire system images available by using clustering managers, a form of "processing virtualization" technology.

Applications or their components are architected to access a cluster manager, and the cluster manager monitors the application and either restarts the application on another system or moves the working application to another system, depending on the type of failure. Workload management and migration are managed at a low level inside the OS.

As with the other types of clusters, this approach relies on data being housed on a separate part of the cluster devoted to storage access or on a SAN, so that it remains available even if the systems hosting the applications themselves fail. Also, access virtualization technology is utilized so user access can be easily migrated from the failing systems to the new systems automatically.

In this case, a form of processing virtualization—cluster and workload management—is the main virtualization technology in this type of cluster. Storage hosts can be housed in the same or different datacenters.

Suppliers such as Citrix, Microsoft and VMware offer this type of processing virtualization today.

Database and Storage Clusters

Another use of the traditional cluster configuration is to support parallel- or grid-oriented databases or storage. The cluster manager's ability to support specially architected database technology, such as Oracle RAC or IBM PureScale DB/2, are typically database offerings designed for this type of configuration. While it does enhance database availability, the primary goals are database performance or scalability. New NoSQL databases, such as those offered by Couchbase, FoundationDB and MongoDB, are also designed to support large-scale clusters.

Continuous processing systems are better hosts for critical functions.

Special-purpose SANs are also built using this type of technology. Often, general-purpose systems access data stored in this system over a special-purpose, high-speed SAN.

Virtual Machine Software Emerges

A couple of processing virtualization technologies, virtual machine (VM) software and OS virtualization and partitioning, have emerged as the focus of today's HA strategies. Entire systems are encapsulated and workload monitoring and management combined with system image migration technology are replacing previous forms of clusters.

Applications running in these system images don't need to be written to use cluster APIs. If a virtual system appears to be in trouble due to a hardware failure, the entire virtual system can be moved to another host. This is a significantly simpler approach to HA. Failover can be managed in seconds or minutes.

Continuous processing systems, however, are better hosts for critical functions. Failover in that type of environment can take place in milliseconds or microseconds.

The Design Center Has Changed

The industry is in the final stages of a significant design center migration. In the past, the design center was keeping systems available and reliable through the use of special-purpose hardware and firmware. Now, the design center is using virtualization technology to assure that applications and their underlying components are available.

The new assumption is that hardware, regardless of whether the hardware is a system, network component or storage component, is going to fail; and properly designed software can provide a low-cost, simple-to-use strategy to address that failure.

Once a system image is encapsulated, it can be hosted on a local system, a system in another datacenter or on a system in a cloud services provider's datacenter.

How Much Availability Do You Really Need?

We're now in a world in which enterprises increasingly need their systems to be constantly available.

We're now in a world in which enterprises increasingly need their systems to be constantly available, and in a world in which these same enterprises need to do the most with a reduced IT budget and staff.

Enterprises would be well advised to review their portfolio of applications to determine how much availability is necessary for each application, rather than how much is available. Some applications cannot be seen to fail, while it may be OK for other applications to become unavailable from time to time.

Business-critical applications are be best hosted on continuous processing systems. Less-critical applications might be happy executing on a cluster or even out in the cloud somewhere.

My advice is select the HA strategy right for each application, rather than using a "one-size-fits-all" approach. 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.
