



Preventable disasters - How come many virtualization users never think availability?

By Dan Kusnetzky, Principal Analyst
Sponsored by Stratus Technologies, Inc.

INTRODUCTION

Many organizations that are embarking on the journey to a more virtualized environment are more concerned with performance, consolidation or optimal use of server resources than preventing disasters. Availability and reliability just isn't part of their thinking until they actually experience a problem.

This paper will examine the following concepts:

- ☒ What is virtualization?
- ☒ Creating a highly available environment
- ☒ Hardware versus software approaches to availability
- ☒ Fault tolerance versus high availability
- ☒ What is enough availability?
- ☒ How can we make an informed decision?

In the end, availability, like management and security is best when it is “baked into” the architecture of an environment rather than something brought into the environment after the architecture is in place and solutions are in place..

WHAT IS VIRTUALIZATION?

Virtualization is a way to create an illusionary environment for IT-based solutions that provides benefits over a physical environment. Applications and their underlying components are “abstracted away” from the hardware supporting them using sophisticated hardware and software. They appear to execute in a “perfect” environment. A logical or virtual view of underlying physical resources is presented to those IT-based solutions. This view may be strikingly different than the physical view.

The goal usually is one of the following: higher levels of performance or scalability, optimal system utilization, scalability, reliability/availability, agility or to create a unified security and management domain. It may also be utilized to allow older, established applications to continue to run even though resources such as the systems, network adapters, storage adapters or storage devices that were a central part of the IT-solution's design may be obsolete or not present in the physical environment.

All of this must be possible without requiring the administrative and operational staff to be aware of the complexities of where these resources reside, whether they're physical or virtual resources and understand where these resources actually reside or how they all fit together.

The Kusnetzky Group has developed a framework or model that describes the layers of technology that can be used together or separately to create a virtual environment. The layers and naming conventions found in the model will be used throughout this paper. The model can be seen in Appendix A.

CREATING A HIGHLY AVAILABLE ENVIRONMENT

There are a number of ways to create a highly available environment. All of them require some level of redundancy. That is, duplicating important resources and orchestrating use of these resources so that work can continue even if something fails. The failure or outage can be planned to allow system maintenance or unplanned. Which resources should be replicated for availability depends upon an organization's requirements for availability, the budget and expertise that's available on staff to plan, deploy and manage the environment.

SOFTWARE APPROACHES

There are different types of software that can increase levels of availability and reliability depending upon if the goal is reliable access, reliable application execution, reliable processing for an entire stack of software, reliable access to storage, or reliable networking. Products from suppliers such as Cassatt, Citrix, HP, IBM, Microsoft, Novell, Scalent Systems, Surgient or VMlogix might be used to create a more highly available environment. As one would expect, each of these types of software can increase the level of availability at a price. Each requires memory, storage, processor power and requires someone set them up, keep them up to date and deal with operational issues. Software failover is not instantaneous. It takes time for the systems to recognize an outage, determine what to do about the outage and then take action to correct the outage. Depending upon the method and technology selected, the failover may take minutes to hours.

HARDWARE APPROACHES

Redundant hardware can be acquired in multiple forms including: installing multiple independent computers, installing a blade computer having multiple processors or installing a system designed from the ground up to be fault tolerant. Suppliers such as Dell, HP, IBM and others would be more than happy to offer a customer multiple systems, blades, power supplies and the like. They would point out that these general purpose systems can be harnessed together to create an environment that is highly available. These systems rely on virtualization technology to detect an outage, determine its cause and execute a strategy that allows workloads to continue to support the organization.

Furthermore, these redundant systems could all be in use hosting workloads (hot backups). Another approach is to have "spare machines" that are idle, waiting to pick up workloads in case of an outage (cold backups). A third approach is to have three or more systems clustered together in the belief that all of them would not fail at the same time. So, workloads could always find a home.

In some cases, a better approach is to use systems that were designed from the ground up to be tolerant of failures and continue running even if something inside of the box fails. In this case, no failure is seen by the workloads hosted on these systems, as these systems eliminate failures rather than minimizing failover and recovery time.

SYSTEMS DON'T RUN IN ISOLATION!

Regardless of which approach is selected to keep systems running, it's important to remember that systems need access to storage containing applications and data as well as access to a reliable network. If these components are not considered in an organization's availability architecture, it may not matter that

systems are still running. If the data is no longer available due to a storage system or network outage, it makes little difference if the systems are running.

FAULT TOLERANCE VERSUS HIGH AVAILABILITY

Of all of the configurations mentioned above, only the last configuration, the fault tolerant system, offers continuous or non-stop computing. All of the others require software to monitor system operations and then execute either a failover or restart operation of a function somewhere else. The fault tolerant system completely hides hardware failures. The workload continues running even though something has failed or has been taken down for routine maintenance.

THE NINES

Often suppliers of availability solutions speak about what level of availability they're offering as an uptime percentage. Let's look at uptime to gain an understanding of what adding a "nine" will do to an organization's exposure to downtime.

Monthly Uptime	Monthly Downtime	Seconds Down per Month	Minutes Down per Month	Hours Down per Month
99.0%	1.0%	26,280.00	438.000	7.300
99.9%	0.1%	2,628.00	43.800	0.730
99.99%	0.01%	262.80	4.380	0.073
99.999%	0.001%	26.28	0.438	0.007
99.9999%	0.0001%	2.63	0.044	0.001

Although workload uptime is often dependent on many factors including the uptime of their systems, storage devices, network and probability of staff error causing a slowdown or failure, the chart above shows that a system supplier offering 99% is really telling its customers that, on average, that they'll experience just over 7 hours of downtime in any given month. While that might be good enough for some workloads, it is woefully lacking for others.

If we add a "nine" to that uptime percentage to provide 99.9% uptime, those same organizations would experience nearly three quarters of an hour of downtime in any given month.

Depending upon the type of system, 95% to 99% uptime is a common range for a single system.

Those offering a clustering-based solution, which, by the way, includes most virtual machine migration-based clusters, would point out that they offer between 99.5% and 99.9% uptime. Even if one of these solutions could offer 99.99% uptime, the organization would experience 4.32 minutes of downtime in any given month. A financial institution could lose millions of dollars if their EFT or trading systems are down that long.

Stratus, a long time supplier of fault tolerant systems, would point out that just isn't good enough for critical applications that can not be allowed to experience **any** downtime. Stratus wants organizations to have six nines, that is 99.9999% uptime. That means experiencing less than 3 seconds of downtime in any given month. Stratus' ftServer systems have operated at or near that mark since first being launched in 2002. (See the Uptime Meter on the Stratus website for hardware and OS uptime calculated daily across its installed base of servers.)

WHAT IS ENOUGH AVAILABILITY?

How much "availability" is enough is a question very much like asking "how long is a piece of string?" Once a specific piece of string is selected, it can be

measured and its length determined. Otherwise, there really isn't a way to answer to the question.

An organization's workload must be evaluated on a function by function basis to determine how long an outage can be experienced for that function before the company suffers ill effects. Some workloads are important but not business critical. Some functions are business critical but not mission critical. Most organizations have a mix of availability requirements. Failure of mission critical systems often leads to the failure of the organization.

It is really important to understand that some functions, such as Email or a collaborative application that have been ancillary in the past, have become business critical as organizations have moved into the 24x7x365 world of the Internet.

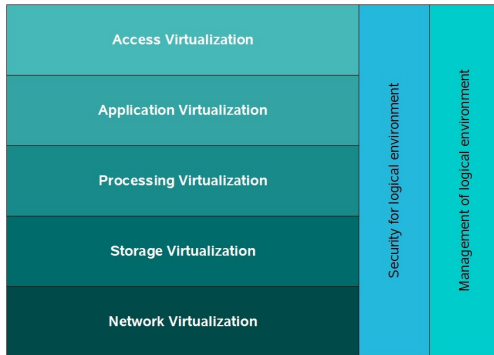
The cost of downtime must be factored in to an organization's thinking

HOW CAN WE MAKE AN INFORMED DECISION?

Making an informed decision about the architecture of important systems has always been challenging. Suppliers of hardware, software, hosting and managed services all want to be a part of an organization's availability solution. It would be wise to seek out suppliers having a long track record of success in the area of fault tolerance and high availability and their partners to obtain the best advice before embarking on the journey to availability and fault tolerance. By planning ahead, many disasters are avoidable.

Appendix A

THE KUSNETZKY GROUP MODEL OF VIRTUALIZATION



Kusnetzky Group © 2007

Access Virtualization — hardware and software technology that allows nearly any device to access any application without either having to know too much about the other. Functions such as terminal services and presentation managers would be found in this layer.

Application Virtualization — software technology allowing applications to run on many different operating systems and hardware platforms. Application virtualization creates a view that all of an organization's industry standard systems hosting the same operating systems as a pool of resources that can be orchestrated dynamically to meet service level objectives, respond to planned or unplanned outages or even to consolidate workloads onto a smaller number of physical systems so that systems may be shut down to reduce power consumption and heat production sys-

tems so that systems may be shut down to reduce power consumption and heat production. It also makes it possible for multiple previously incompatible applications or several versions of a single application to run simultaneously on the same physical system.

Processing Virtualization — hardware and software technology that hides physical hardware configuration from system services, operating systems or applications. This type of Virtualization technology ranges from the ability to make one physical system appear to be many or many systems appear to be a single computing resource. Usually this type of technology is deployed to achieve goals ranging from high levels of performance, scalability, reliability/availability, agility or consolidation of multiple environments onto a single system. Different types of processing virtualization are needed to achieve these goals. This layer of virtualization technology supports grid computing, single system image clustering, HA/failover clustering, client and server virtualization as well as operating system partitioning/virtualization.

Storage Virtualization — hardware and software technology that hides where storage systems are and what type of device is actually supporting applications and data. Storage virtualization makes it possible for many different physical systems to share a single storage resource to reduce the costs of purchasing physical storage for each system or for this storage resource to be replicated several times in different datacenters to facilitate disaster recovery.

Network Virtualization — hardware and software technology that presents a view of the network that differs from the physical view. Network virtualization makes it possible for many workloads to share the same network environment securely. Clients may be allowed to only see servers they're allowed to access. Servers may only see clients they support.

Management of virtualized environments and security — Two of the most important layers of virtualization are the layers that manage and make secure all of the other layers of virtualization technology. This software technology makes it possible for multiple systems to be provisioned and managed as if they were a single computing resource. Without this layer of technology, organizations would face greater complexity and costs in a virtual environment than they did when they were using only physical systems.

RELYING ON PERFECTION

Unfortunately, a good deal of the layers of technology that make up a virtualized environment have been designed as if the underlying hardware and administrative processes are absolutely perfect, never experiencing outages. This, of course, is ridiculous on the face of it. Hardware fails. IT administrative and operational staff make mistakes that result in slow downs or outright failures. Architects of virtualized environments should take availability into account when building an organization's virtual world, but often do not.