



Building a highly resilient IT infrastructure:

*Increasing availability using
next-generation data center
technologies to achieve 100% uptime*

Executive Summary

There is a heightened initiative throughout the IT industry to build highly available infrastructures. Today, more and more companies are able to reach the goal of 100% uptime. This white paper outlines the basic tenets of business continuance, discusses the technology enablers available today, and examines some reference architectures for achieving uptime goals. There are so many variables in building a highly available infrastructure, it is important to deeply understand what services the business relies on most and the different ways to protect those services.

Disaster Recovery and Business Continuance Overview

Disaster recovery (DR) and business continuance (BC) continue to be a top priority for organizations. Business continuance is the processes and procedures an organization puts in place to ensure that essential functions can continue during and after a disaster. Business continuance planning seeks to prevent interruption of mission-critical services and to reestablish full functioning as swiftly and smoothly as possible. Technology advancements over the last few years have increased the availability of mission-critical applications that support core business needs while eliminating planned downtime and minimizing the impact of unplanned events.

The primary issue, however, remains: *Is there such a thing as a 100% uptime environment today?*

While there are many elements to consider, the short answer in almost all environments is “Yes.” It is possible to achieve 100% uptime when responding to operational recovery situations. However, disaster recovery efforts are not yet able to achieve this goal—almost always due to the application platforms a company needs to support and its ability to take advantage of infrastructure technologies (e.g., applications cannot be active/active between data centers). That being said, by using virtualization and by understanding where the

industry is moving, a company can, even now, build an infrastructure groundwork able to provide 100% uptime for operational recovery and very close to 99.999% uptime for disaster recovery.

Operational vs. disaster recovery It is vital to understand the fundamental differences between operational and disaster recovery. *Operational recoveries* are typically performed within the same data center as production and mitigate against events such as data corruption, local array failures, and human error. *Disaster recovery*, on the other hand, typically occurs after an outage caused by an external catastrophe, such as fires, floods, and storms. Disaster recovery requires a secondary location and data replication across a WAN infrastructure. Each recovery option has specific events and risks associated with it, as well as physical limitations that determine what recovery time objectives (RTO) and recovery point objectives (RPO) are achievable. These factors will be explored in depth later in this document.

Definitions

RECOVERY TIME OBJECTIVE (RTO): the amount of time required to recover services in the event of an outage

RECOVERY POINT OBJECTIVE (RPO): the amount of data that is lost through recovery efforts

Define your service catalog It is very important when beginning any disaster recovery / business continuity / high-availability project to take three specific steps:

1. Define the IT services to protect.
2. Outline the risks to which each service is susceptible.
3. Create a service catalog that identifies what levels of protection each service will receive.

It is rare that every IT service will receive the same levels of protection in a given data center. Discussions with each application owner determine the business impact if the service were not in operation. Typically these discussions occur when an organization performs a business impact analysis to understand what IT services the business relies on and what will happen to the organization during an outage. That information will provide the requirements the infrastructure team needs to build to.

Understanding uptime When considering RTO and RPO, we must understand the differences between commonly used uptime numbers and evaluate the real effect they have on the business applications.

Figure 1 and Figure 2 show the top eight common uptime calculations in both minutes (Figure 1) and hours (Figure 2). The difference between 100% uptime and 99.999% (i.e., Five 9s) is 5.26 minutes per year; the difference between 100% uptime and 99.99% (i.e., Four 9s) is 52.6 minutes per year.

When embarking on a project to build a highly available infrastructure, we always ask the business “what is your expected RTO and RPO?” The answer we receive is usually measured in hours. A four-hour RPO/RTO mandate equates to an uptime percentage of about 99.95%. Most next-generation availability technologies can provide RTO and RPO in the scope of minutes, very close to a Five 9s architecture, while still being flexible and cost-effective. The goal should be to implement a tiered-recovery methodology that provides the organization with multiple levels of protection and recovery based on the IT services catalog defined above.

Risk Planning and Mitigation

When trying to design and operate an IT infrastructure with 100% uptime, a number of challenges could keep you from reaching your goals. It is important to understand the threats to your data center and how to architect

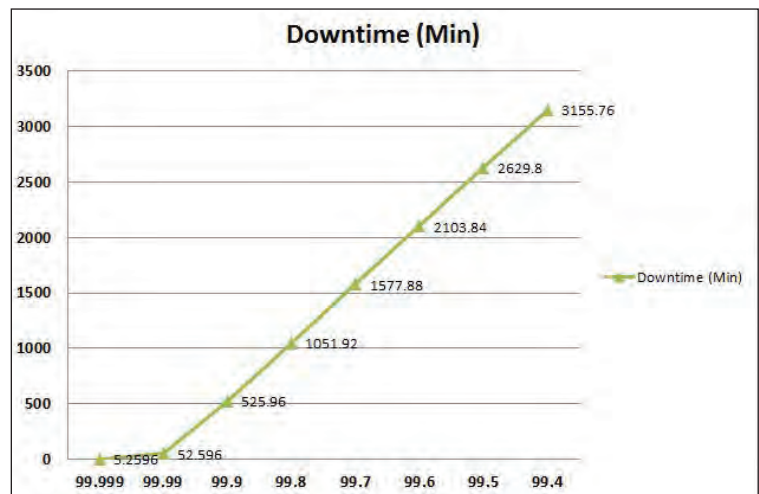


Figure 1—Downtime in minutes




Figure 2—Downtime in hours

100%-Uptime Challenges

OPERATIONAL

- Planned outages
- Data corruption event
- Infrastructure failure

DISASTER RECOVERY

- Natural disaster
- Man-made disaster

solutions that mitigate them. These threats are categorized as either operational recovery challenges or disaster recovery challenges.

Operational recovery challenges These are typically environmental issues within the data center facility: human error, hardware failure, virus attack, or scheduled downtime for upgrades/patching. The key to building a highly available inclusive operational recovery plan is to implement practices and technologies to mitigate as many of the risks as possible, and then protect the environment from events that do not cause immediate downtime. An example would be using technologies that do not require downtime to roll out patches to applications and operating systems within your environment.

Disaster recovery challenges Disasters are more rare than operational outages, but have far greater impact to a data center when they do occur. Disasters are classified into two broad categories: natural and man-made. Natural disasters would include floods or hurricanes, whereas a man-made disaster would be similar to the September 11 attacks or an EMP (electromagnetic pulse) attack in a metropolitan area. When planning for these events, the assumption is that the production data center is completely non-functional and failover to a secondary location is the only way to resume IT services and the business functions they provide.

Build a risk-mitigation plan

Building a risk-mitigation plan is the first step towards designing a resilient IT infrastructure. This also provides the groundwork for a tiered recovery model that can be communicated to the business owners. Laying out potential risks and how the IT department would recover from them allows the business owners to understand the capabilities and be clear on the appropriate response in a given situation. A tiered recovery plan outlines each IT service and the potential risks, with a given RTO/RPO for each risk. This plays a role when designing the architecture and implementing technologies to meet the requirements of the organization. A sample of a basic tiered recovery plan is in Figure 3.

Type of Recovery	Uptime	Technology Solution	Description
Operational Recovery			
Planned Downtime			
OS Patching	99.999%	MSCS Clustering	Implementation of Microsoft Cluster Services would allow patching the operating system of a single node, failing the cluster over that node, and then performing the patch on the secondary system.
Server HW Upgrade	100.000%	VMware vMotion	vMotion allows moving a running virtual machine from one host to another. To upgrade any aspect of one host, put it in maintenance mode. All the servers on that host will migrate with no downtime to other hosts.
Storage HW Upgrade	100.000%	EMC SP Failover	EMC allows hot-add storage and the ability to fail LUNs over to the secondary controller to update or install new storage firmware or software features.
Network HW Upgrade	99.999%	Cisco Supervisors	Cisco Supervisors allow updating code on one supervisor (without impacting network traffic) and then failing over to that node.
Unplanned Downtime			
Virus Attack/ Data Corruption	99.999%	RecoverPoint CDP	In the event of a virus attack or data corruption, use the journal function of RecoverPoint CDP to roll production data back to a point where the data was clean.
Server HW Failure	100.000%	VMware Fault Tolerance	Using VMware Fault Tolerance, a running server can use a shadow of itself on another host. If the primary host has a hardware failure, the secondary host will take over instantly with no downtime.
Storage HW Failure	99.999%	RecoverPoint CDP	In the event of a primary array failure, RecoverPoint CDP would fail over to a secondary array. This would allow resumption of services on the secondary array within minutes with very little data loss.
Disaster Recovery			
Planned Downtime			
Planned Utility Maintenance	99.999%	RecoverPoint CRR with VMware SRM	Use RecoverPoint CRR between data centers to failover to the secondary location in the event of a known utilities issue in the production data center.
Unplanned Downtime			
Tornado/Fire/ Flood	99.990%	RecoverPoint CRR with VMware SRM	In the event of a natural disaster that would take the production facility offline, use RecoverPoint CRR to provide failover to the secondary location.
Terrorist/ Pandemic	99.990%	RecoverPoint CRR with VMware SRM	In the event of a terrorist attack or pandemic outbreak, a fully automated recovery plan would be needed. A third party would execute the recovery plan.

Figure 3—Sample tiered recovery plan

The sample plan shows that each outage type has a slightly different uptime estimation due to the physical limitations of the technologies used to mitigate the associated risks. This is customized for each environment and varies, depending on the infrastructure technologies that are being employed within the data center.

Communicate the plan

Most organizations have DR/BC committees consisting of application owners, IT support personnel, and IT architects. To ensure a successful DR/BC project, it is imperative to communicate the risk mitigation plan and service catalog to the business owners. This ensures they have the proper expectations and, in the event of an outage, they understand the timelines they are dealing with and what steps are being taken to recover the environment. This communication should include the capabilities of the environment and also the limitations.

Building a Highly Available Infrastructure

You've built a services catalog and understand the impact each IT service has on the business in the event of an outage. You understand the risks that could impact uptime. Now it's time to architect an infrastructure using technologies that can mitigate those risks.

Physical vs. virtual environments

A virtualized infrastructure brings many advantages—ranging from power and cooling to lowering costs per compute resource. The main reason to virtualize, though, is to enable a superior level of flexibility, resiliency, and availability in the production infrastructure.

Abstracting the server hardware from the production operating system (OS) through a hypervisor eliminates the need to have matching hardware at both locations to provide disaster recovery. In the physical world, successfully backing up and recovering

production servers requires one of two things: either 1) have the exact same hardware in a DR location or 2) implement a complex system to perform a function called “bare metal restore.” This increases the time to recover those servers and increases the complexity of the environment. In a virtualized world, the hardware is completely abstracted. So long as the hypervisor is similar, the hardware platform can be mixed and matched as needed.

Another powerful advantage of virtualization over physical is the enhanced availability features developed over the last few years. Today, the most highly available infrastructures are virtualized and can take advantage of clustering, fault tolerance, and virtual machine teleportation. These features—highlighted in the next section—have ensured that 100% virtualization is the first step towards implementing any successful DR/BC project.

Clustering (active/active vs. active/passive)

One of the most important availability technologies that has allowed the achievement of such high uptime numbers is the use of clustering technologies. In essence, a clustering technology is used to mitigate the loss of a single server in the event of an outage—most often a server hardware failure or OS issue. The most common form of cluster today is an active/passive cluster of nodes. In an active/passive architecture there is a primary node and a secondary node for a given application resource. If the primary node fails, the secondary node detects the failure and takes over the resources required to run the application. Downtime is usually anywhere from seconds to minutes during this failover, which is typically completely automated.

The challenge to achieving a true active/active environment lies in the database layer of the data center infrastructure. There is currently only one active/active database technology used in some organizations—Oracle RAC. Most organizations tend to have other database implementations that are unable to be built in a true active/active setup. In an active/active environment, multiple nodes can

handle an incoming application request and all nodes have access to shared resources. If a single node fails, there is no downtime because another node simply takes the request and fulfills it.

In the past, most clustering technologies mitigated against operational recovery situations. However, some have been implemented in what is commonly referred to as a “geo-cluster” or geographically dispersed cluster architecture. The concept here is to provide automated active/passive clustering between two data centers in the event of an outage in the primary site. The most common geo-cluster implementations today are in IBM mainframe and AS400 environments, where the entire application is run in a single environment. It becomes more of a challenge in a Microsoft Windows environment and a three-tiered web application infrastructure—Web, app, and database. In this environment it is difficult to synchronize the failover of all the application dependencies to a secondary site and bring them up in an automated fashion. This is why most implementations of disaster recovery technologies require the organization to determine whether or not it wants to fail to the secondary site or if the recovery will happen locally. The decision normally has to do with the type of outage that has occurred in the production facility.

Building the best solution you can today

There is a reference architecture that shows the power of virtualization and brings all the technologies together in the best infrastructure possible at this time. The culmination of years of planning and development, it provides the strongest multi-tiered recovery implementation we have encountered. In the following architecture, you will see that through the use of a few key infrastructure technologies this organization mitigated the majority of its infrastructure risks and provided the maximum uptime possible in its environment—all while optimizing costs and management.

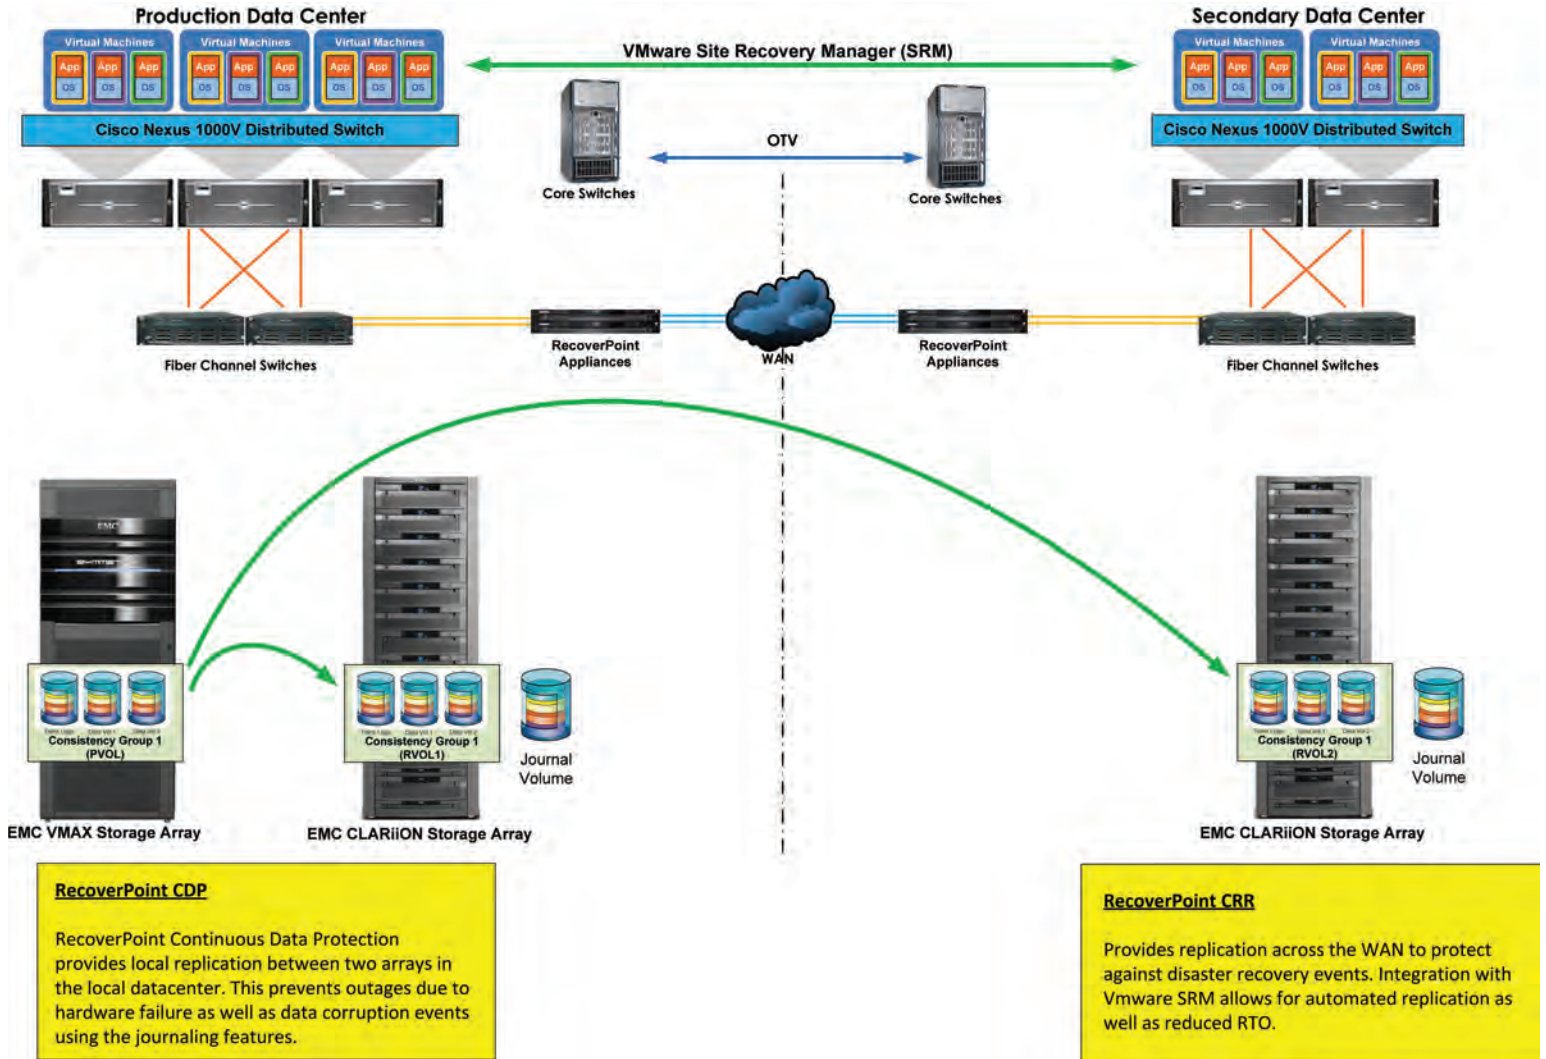


Figure 4—Sample reference architecture

Conclusion

There continues to be a substantial amount of research and development in the area of disaster recovery and business continuity technology infrastructure. Limitations that were once deemed impossible to overcome have been removed and problems continue to be solved, promising geographically dispersed active/active data centers to come.

While these technologies are still being developed, by following the reference architecture outlined here it is possible to build a 100% operational uptime environment now. In the near future, companies will be able to provide 100% uptime, even in the event of a full disaster recovery scenario. The key to achieving this is to follow a path of virtualization—implementing technologies today that are compatible with the future road maps.

Disaster recovery and business continuance can be a daunting proposition, but Ahead has built a proven methodology to accelerate the process and ensure a successful deployment—using highly trained professional services staff and an in-depth knowledge of desktop and server virtualization.

You're not alone. Contact Ahead to find out how you can get started. (Call 312.329.7880, email sales@ThinkAheadIT.com, or visit www.ThinkAheadIT.com.)



About Ahead

Ahead is a leading provider of next-generation data center solutions, with specific expertise in virtualization and cloud computing architectures. Ahead designs and implements agile, service-oriented architectures to help clients transform their data centers from a technology-focused environment to a service-delivery platform that brings agility to the business and drives out cost. Solutions are built using a unique, simple, scalable, and repeatable methodology that speeds delivery and lowers risk.

*Ahead, purpose built to drive private cloud enablement, employs a unique methodology called **THINK | LOOK | PLAN | MOVE**.*

*In the **THINK** phase Ahead works with clients to identify challenges and to understand and define current business objectives and strategies.*

*In the **LOOK** phase Ahead identifies all service elements and dependencies, performs gap analysis, and evaluates appropriate technologies.*

*In the **PLAN** phase Ahead develops architectures and TCOs, along with tactical and strategic implementation plans.*

*Finally, during the **MOVE** phase Ahead initiates its plan and executes on delivery of its offerings.*

For more details, visit www.ThinkAheadIT.com.

The material in this document is the proprietary information of Ahead, LLC.

All products, trademarks, and copyrights herein are the property of their respective owners.

©2011 Ahead, LLC. All rights reserved. revised 2/21/11