# Five Things You Need to Know Today about
# Disaster Recovery Planning

NetIQ®

# Five Things You Need to Know Today about Disaster Recovery Planning

## Table of Contents

# Introduction

As a busy data center manager, you might be struggling to keep up with the challenge of continually doing more with a smaller IT staff and budget. IT has become an integral part of business operations—in fact, most, if not all, of the infrastructure you support is critical for keeping your company up and running—but doing more with less has most likely forced you to put some projects on the back burner.

Chances are good that disaster recovery is one of the projects you'd like to move to the front, and for good reason. It's nearly impossible to follow the news without feeling bombarded by stories that bring home the potential risks out in the world. Whether it is turbulent weather, natural disasters or man-made accidents, something bad is always happening.

When was the last time you looked at your disaster recovery plan? What do you expect would happen if the ceiling collapsed in your data center?

What would you do if an employee forgot to unplug a humidifier in his or her cubicle and the power grid feeding your servers and storage imploded?

This guide will cover five topics you need to understand about disaster recovery in today's data centers.

1. **Multi Platform** – Mixed platforms are everywhere (the new face of IT)
2. **Virtuality** – Virtualization changes everything
3. **The Next Big Thing** – Cloud computing is the new model of IT service delivery
4. **Cost Savings Galore** – Measuring returns on disaster-recovery investments
5. **The Hidden Secret** – Disaster require planning and testing for higher confidence

Many of these require new thinking when it comes to disaster-recovery planning. Soon, you'll be ready to tackle the new realities of disaster-recovery planning, with new approaches to protection that can make your life easier.

# Multiple Flavors Aren't Just for Ice Cream Anymore

# 1 Multi **Platform**

Multi-platform data centers are no longer an exception. They have become the rule. Several years ago, most data centers were quite homogeneous; all of the data residing there typically came from one vendor's devices. This made things relatively easy to manage. Data center managers had only one set of instructions to learn and one phone number to call when something went wrong. The problem was this: The data center devices were called mainframes, and their price tags were as massive as they were.

As the IT world continued to grow, several engineers and innovative companies started to design much cheaper processors, allowing other innovators to build real servers using PC components. These new servers used a new set of common standards, and they had a much lower cost barrier—even small organizations could leap it and enter the computer-generated data game. Large organizations took advantage as well, and both selected products from a vendor-rich buffet of x86 building blocks to construct huge data centers. Some call this time the golden age of the data center. Analysts simply called it the client-server model.

The move from mainframes to smaller servers made the data center strategy simple: When you needed another application, you just bought another server. Servers were cheap, so why not? The system worked well for a time, until all those servers began to take up way too much space and use way too much power. And most of the servers were mostly underused most of the time.

The solution to these problems was supposed to be virtualization. Virtualization allowed data centers to consolidate computing resources, shrinking the physical computing space and its energy consumption to something more manageable. But all of these data center changes overlapped. Nothing ever went away completely. Data centers just got a little bit smaller, and energy bills continued to grow.

As a result, today's data center is typically a hodgepodge of platform and vendor flavors: a bit of mainframe, a little x86 and a whole lot of virtualization.

So what happens if the power goes out? Unfortunately, the hodgepodge of solutions in the data center leads to a recovery plan that's a bit of this and a little of that, and so forth. In other words, you need one plan for the mainframe, one for the Linux* servers, one for the Windows* servers and probably a virtual recovery plan, too.

**More flavors, more confusion, more headaches.**

# 2 Virtuality

The typical process for recovering physical servers is:

1. Find or buy equivalent or compatible physical servers.

2. Install operating systems on them.

3. Install patches and updates.

4. Install applications.

5. Install application patches and updates.

6. Upload backup data.

Using virtual recovery, you can make almost all these steps things of the past. The process to recover a physical server can be as simple as:

1. Power on its virtual equivalent.

2. Be thankful you no longer have to deal with physical servers.

This far more elegant approach can save a great deal of time and effort.

## Virtual This, Virtual That—Let's Get Real

Hidden within your data center's jumble of mixed platforms is a secret: You can use all those underused virtual resources of yours to recover more than just identical virtual machines. One wonderful thing about virtualization is that you can use virtual machines to run many different types of workloads. You can, and should, apply this flexibility to disaster recovery. Why not create a virtual recovery platform that can offer protection for all your workloads, whether they are running on physical or virtual servers, or on Windows or Linux operating systems?

Virtual recovery plans can simplify, and in some cases eliminate, many of the platform-specific headaches of recovering from a disaster.

# When Did Weather Metaphors Become So Common?

Unless you've been living under a rock, you've probably heard that cloud computing is the next big thing. You might have even seen a TV advertisement or two talking about cloud computing. What is it, and can it affect, or better yet, improve disaster recovery?

To answer the first question with a definition, some might say that the term cloud computing is synonymous with the term buzzword. But cloud computing actually has a number of well-accepted characteristics that define the term.

A cloud-based infrastructure looks like this:

1. Shared infrastructure (or, plays well with others)

2. Rapid elasticity

3. On-demand, self-service resource-acquisition model

4. Per-use billing

5. Standard Internet protocols

6. Measurable usage

If you think about this list, you'll realize that disaster recovery is a natural benefit of using the cloud-based resource-consumption model. Disaster recovery is all about the unknown—not knowing when you will need extra resources, how many resources you will need or for how long you will need them. The on-demand, elastic model of cloud computing is perfect for these rather "cloudy" requirements, don't you think?

But how do you get to the cloud? The first step on the stairway to the cloud is virtualization. After all, virtualization is the main underlying technology that makes the cloud-delivery model possible.

So if you aspire to a silver-lined disaster-recovery future, you should virtualize as many of the pieces that make up your disaster recovery plan and computing infrastructure as possible.

3

## The Next Big Thing

# Wait…This Sounds Expensive

Virtualization has already proven its cost-saving benefits in production data centers around the world. By reducing the physical server sprawl created by long years of accumulating commodity x86 servers, virtual machines have proven themselves to be the cheaper way to run workloads. Not only do you dramatically lower your server hardware costs, but you also lower the cost of powering, cooling and maintaining server hardware. However, if yours is like most data centers, many barriers stand in the way of adopting virtualization as the universal platform, despite its cost benefits. While nearly everyone agrees that it is a cheaper way to run workloads, many data center managers are concerned about virtualization from the performance and security standpoints.

Disaster recovery faces far less scrutiny than do workloads when it comes to infrastructure changes. This makes disaster recovery an area where you could immediately adopt a fully virtual strategy. Since disaster recovery is an often-forgotten element of your overall IT strategy, implementing a better performing recovery infrastructure will be welcome by everyone involved—or it might go unnoticed. (This is a good thing, by the way. Getting everyone to notice your brilliant new disaster recovery plan sometimes takes an actual disaster.)

But even if better performing recovery capabilities go unnoticed, the immediate cost savings won't. As part of a disaster recovery plan, one virtual host server can take the place of up to twenty or more physical standby servers. Think about it. The expensive duplicate infrastructure that your data center once needed can now be a thing of the past. Eliminate the burden of keeping backup versions—make, model and vendor duplicates—of all the servers you run. Replace all of them with a simple pool of virtual resources. This drastic reduction in infrastructure immediately reduces the cost of your disaster recovery plan.

But infrastructure isn't the only thing that costs money. Your time is ultimately more valuable than the gizmos you manage. Virtualization also reduces the overhead and labor associated with older backup-and-recovery protection approaches.

## Why bother with convoluted, multi-step recovery processes when you can streamline all your administrative tasks with virtualization?

With virtual machines, everything from day-to-day maintenance to recovery, and even to testing, can be as simple as a few clicks of your mouse.

By implementing virtual disaster recovery, you can effectively get a double bonus in real costs savings. Cheaper to buy, and cheaper to run—what could be better?

## 4 Cost Savings Galore

# Hope Is Not a Strategy:
## The Hidden Secret of Disaster Recovery Planning

For many years, the hidden secret of disaster recovery has been testing. Are you surprised? Think about it for a minute. When was the last time you tested your ability to restore a workload? How many did you test? And (here's the kicker) what fraction of your total workloads was that? Backing up and restoring data is a real pain. Getting everyone involved to agree on a testing schedule is hard enough, let alone having to rely on antiquated technologies like tape backups once you to get there. When the difficulties associated with disaster-recovery testing seem impossible to overcome, it simply never gets done.

No or infrequent testing kills confidence. After all, how can you be sure your plans will work if you don't try them out once in a while? Even worse, in today's business world, everyone expects you to publish service levels and guarantee them; how can you do this without an accurate prediction of what a worst-case scenario would look like? Users now demand accurate expectations. "One or two days," or other ambiguous service levels are no longer acceptable. How can you publish an accurate number if you never run tests to prove it?

Again, virtualization can be the technology that magically overcomes the difficulties preventing effective disaster-recovery testing. Virtualization eliminates the problems associated with a bare-metal workload restoration. You don't need to match hardware, or go through multiple steps just to get a test server up and running. This easy management allows you to simply select the virtual machines you want to test, create copies of them and power them on. Virtualization allows you to implement a testing process that is completely non-disruptive to your production processes.

Safe and easy—what more can you ask? Easy testing can provide assurance that your data-recovery plans are effective and enable you to confidently set service levels.

With the ability to run disaster-recovery tests easily and often, you can accurately measure the time you expect workload recovery to take. Instead of hoping that you can recover from an outage in one or two days, you can now guarantee and even publish an accurate recovery time objective (RTO).

You can even use your guaranteed RTO as a competitive advantage for your company. Everyone loves routine, so people want to know that in the event of a disaster, you can help them get back to their routines as quickly as possible. The ability to publish service levels you can confidently meet is tremendously valuable for your—or any—organization.

5 The Hidden Secret

# Fantastic! What's Next?

Take time to review your current disaster-recovery plans. This new virtual world makes disaster recovery easy, but you can't move to it without some effort. Following are a few suggestions about where to start.

Arm yourself with as much knowledge as you can. Try initiating conversations with users in your organization. Ask people how important their jobs are (everyone loves this question). Their answers will help you build a better understanding of what your RTO needs really are for each workload running in the data center.

Use the simple uptime-downtime chart that precedes this section to talk to your users about the dollar-cost impact to which these service levels might translate. If your order processing application, for example, demands three nines (0.999, or 99.9 percent) of availability, how many transactions do you expect would be lost or delayed, and what would be the cost of these lost or delayed transactions if you allowed for 24 hours of annual downtime? Does this cost align with the cost of providing a three-nines guarantee?

Remember, disaster recovery isn't quite what it used to be. With mixed platforms, virtualization clouds, cost confusion and testing, you have an awful lot to consider. Hopefully this guide has given you some food for thought as you consider updating or changing how your organization looks at disaster recovery.

## Allowable Downtime for Workloads

| Availability SLA | Required uptime hours-per-year | Allowable downtime-per-year |
|---|---|---|
| 90% (0.9) | 7884 | 36.5 days |
| 99% (0.99) | 8672 | 3.6 days |
| 99.9% (0.999) | 8751 | 8.7 hours |
| 99.99% (0.9999) | 8759 | 52 minutes |
| 99.999% (0.99999) | 8760 | 5 minutes |

# Arm yourself with as much **knowledge** as you can find and start having conversations about **disaster recovery** with people in **your organization**.

Here are some additional disaster recovery resources from NetIQ:

**Product information:**
http://www.netiq.com/products/forge/index.asp
http://www.netiq.com/products/protect/index.asp

**Customer success stories:**
https://www.netiq.com/success/stories/cash-financial-services-group.html
https://www.netiq.com/success/stories/metropolitan.html
https://www.netiq.com/success/stories/italian-central-government.html
https://www.netiq.com/success/stories/euler-hermes.html

## About NetIQ

NetIQ is a global, enterprise IT software company with relentless focus on customer success. Customers and partners choose NetIQ to cost-effectively tackle information protection challenges and manage the complexity of dynamic, highly-distributed business applications.

Our portfolio includes scalable, automated solutions for Identity, Security and, Governance and IT Operations Management that help organizations securely deliver, measure, and manage computing services across physical, virtual, and cloud computing environments. These solutions and our practical, customer-focused approach to solving persistent IT challenges ensure organizations are able to reduce cost, complexity and risk.

To learn more about NetIQ's industry-acclaimed software solutions, visit **www.netiq.com**.

For a complete list of our offices in North America, Europe, the Middle East, Africa, Asia-Pacific and Latin America, please visit www.netiq.com/contacts.

NetIQ®