Your Managed Cloud Hosting Provider That Cares

Backup and Disaster Recovery Planning Guide for CIOs

HANDY Networks LLC

Backup and Disaster Recovery Planning Guide for CIOs

Handy Networks LLC
1801 California Street
Suite 240
Denver, Colorado 80202

(303) 414-6910

http://www.HandyNetworks.com

Schedule your complimentary Data Backup and Disaster Recovery Consultation now at:

http://content.handynetworks.com/dbdr

Introduction

As a CIO or IT director, you have a lot on your plate. In addition to leading the software engineering and systems administration teams, you're always looking to improve and develop your teams and engineering processes. Your company's strong growth has put a lot of pressure on most parts of your teams and the internal customers they support.

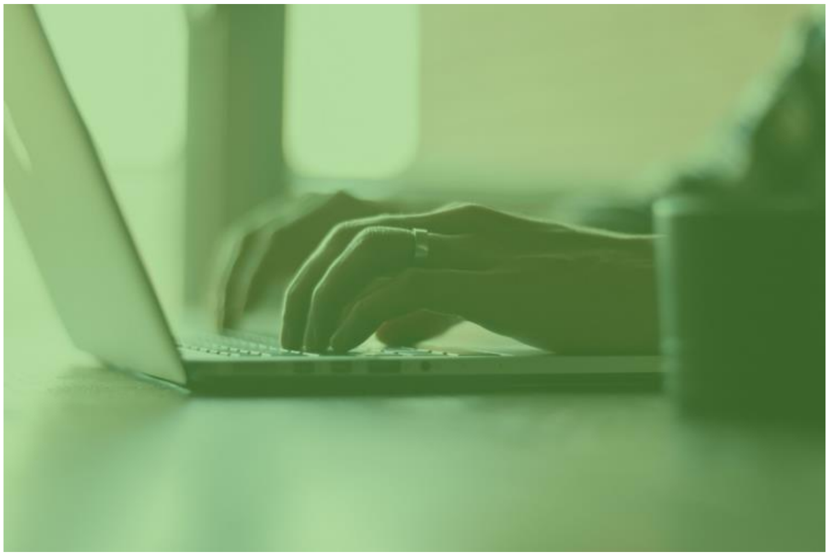

You face a lot of different, interrelated technology and staffing challenges. At the same time, your company is often criticized for not having up-to-date software – which often leads to website and application crashes, and makes it tough to provide timely support.

Your understaffed team struggles to have continuity in its engineering development, which often leads to extra stress and extra work with little advanced notice. However, even with all of this going on, your company is still depending on you and your team to have a rock-solid data backup and disaster recovery strategy. In this eBook, you'll learn what it takes to protect your company's digital assets and overall business continuity.

Data Backups

Like insurance and risk management, data backups are one of those business areas that most people pay very little attention to until it's too late. However, as your company's CIO or IT director, you can't afford to indulge in data backup denial and procrastination. There's simply too much on the line.

But at the same time, you don't have the luxury of devoting the majority of your time and budget to data backups. However, you do need to prioritize a few basics.

The Best Data Protection Systems on the Market Today

While there are potentially dozens of data backup options you could consider, we've done the hard work for you, vetted the contenders, and have a few shortlist options to share. We use a number of different data protection platforms and products. We're especially big fans of Veeam and R1Soft CDP.

Veeam

Veeam is interesting because it was designed initially just to handle virtualized VMware environments and do backup and recovery just for virtualized VMWare environments.

Over the years, Veeam has expanded its capabilities. Now Veeam can backup Hyper-V and use agents to do bare-metal backups of bare-metal Windows and Linux servers. Veeam can also be used as part of a disaster recovery strategy.

Veeam has a lot of beneficial features that you can leverage:

- Veeam Cloud Connect replicates your backup onsite (if you are hosted in the cloud) or offsite (if you are hosted on-prem).
- Veeam can spin up virtual machines instantly (depending on the infrastructure you're using), so if you need to recover quickly, you can spin up the VM. The performance is not necessarily going to be as good as it would be if it were on your primary environment.
- Veeam supports stateful backups of Active Directory, Exchange Server, SQL Server, and Oracle.
- Veeam backup/restore is solid, easy to perform, and it has a number of option to leverage SAN based snapshots for backups, further improving backup performance.
- Veeam has the ability to use Common Internet File Systems (CIFS) for the storage backend, making it easier to scale storage on a per client basis.

As robust of a data backup platform as Veeam is, it isn't for everyone. Why? Veeam comes with a pretty hefty price tag. The price for Veeam works out to about $12 a month per VM, just from a licensing perspective. If you have 100 VMs, you're paying $1,200 a month just for the software licenses.

For some organizations, there are no issues with these costs. Other clients are much more budget-conscious. Veeam also lacks an intuitive, web-based self-service portal that you or your customers can rely upon to do their own restores. For our clients that require a more affordable price point and self-service portal, we recommend R1Soft CDP.

R1Soft CDP

As one of their initial customers on their R1Soft CDP's 2.0 release many years ago, it's been interesting to see the product evolve.   At this point, we are pleased with R1Soft's feature set, especially relative to its price point.

R1Soft was initially conceptualized as a backup system to be used by data centers and hosting companies and has a number of features that facilitate self-service restores for end users, through a web-based interface and control panel integrations.

R1Soft CDP has some things Veeam doesn't:

- ✓ **Nice web interface**
- ✓ **Easy-to-apply policies from a centralized perspective**
- ✓ **Multi-tenant capability**
- ✓ **Integrations with cPanel for end-user hosting restores**
- ✓ **Runs on Linux which saves on licensing costs**

With this feature set, we can provide our users and customers with logins to the backup systems, so they can monitor and restore their own files if they want to.   R1Soft CDP is also very inexpensive. We have a very aggressive licensing deal with R1Soft that's much cheaper than Veeam.

As a result, R1Soft CDP is usually the default offering that we bundle with our bare-metal dedicated hosting environment and our entry-level private cloud hosting.

Autotask

Another data protection product we offer, and use daily internally, is Autotask File Sync and Share. Autotask is not strictly a backup system. Rather, think of Autotask File Sync and Share as enterprise-level Google Drive. It has a lot of features:

- Single sign-on
- User groups and projects
- Shareable links
- Password protection
- Encryption
- PCI compliance
- HIPAA compliance
- 180-day file version history
- Basic mobile device management (including remote wipe for laptops, tablets, and phones)
- Configurable file/folder level backups of data on your local PC.

One of the beautiful things about Autotask is it keeps revisions of all files for 180 days. If a non-technical end user makes a change to a file realizes they need to revert to a previous revision, the end user can login to the web interface, show the file revision, and download it directly.

Typical file restore requests require a number of interactions between the end user and the IT helpdesk, to determine exactly the file that needs to be restored, the date it needs to be restored from, and so on.

With AutoTask FSS, end users within your organization can handle this task completely on their own. This feature alone is a huge productivity win for both IT staff and non-technical staff.

Autotask is also a great way to protect yourself from ransomware on your file servers. Let’s say you get an outbreak of some sort of cryptolocker malware on your network and it encrypts all of your file shares.

With Autotask, you can just go in and revert everything back to yesterday’s copy--before the malware broke out--and your data is almost instantly restored. You can mitigate the malware on your network, and you don’t have to payout a massive amount of money in bitcoins to these hackers.

SAN Snapshots & Replication

SAN based snapshots and replication are also another important element of a data protection plan. All enterprise-class storage platforms, such as Nimble Storage, EMC, NetApp, Compellent, all support snapshots of some sort or another. A snapshot is simply a point in time copy of all of the data associated with a particular SAN volume (or LUN.)

You might think this takes up a huge amount of space. Fortunately, modern SANs rely upon block level deltas, which means your snapshot only consumes the amount of blocks that changed from the previous snapshot.

We use SAN snapshots for a number of uses, including dev/test/troubleshooting. However, we also use them as part of our data protection strategy. For instance, we backup all of the Hyper-V VMs running in our Highly Available IaaS platform using R1Soft CDP.

Yet, we still take SAN snapshots of the same data to facilitate quick and painless restores of large data sets. SAN snapshots also make it easy to revert to a previously known-good set of data when facing unexplained problems.

SAN replication offers an extra layer of data protection. We use SAN replication to create block-level copies of critical internal and customer data to remote data center locations. We can then couple these SAN level snapshots with Microsoft and VMWare tools to create simple and effective disaster and recovery solutions for our clients.

Making the Business Case for Proper Data Backup Investment Levels

How can CIOs and IT Directors make the business case for investing in data backups at the right level? That's a hard question to answer because it's highly context specific.

In a perfect world, the CIO is the one who's making up the business requirements for whatever data protection requirements the company needs to have. These data protection requirements are driven by:

- The industry that the company is in (and the regulatory requirements that go along with that industry)
- The size and maturity of the organization
- Tolerance for risk
- Budgetary realities

Our data protection options start at $120 per year, per server being backed up. With such an affordable price point, there's no reason your company should not be able to afford at least basic R1Soft CDP based backups.

Sidestepping Common Data Backup Problems

What are the biggest problems with most data backup systems? One thing we've run into--which Veeam and R1Soft CDP both address--is older backup technology that has some odd things to overcome. This can be especially challenging if you're looking to use a backup-to-disk methodology.

A lot of the older backup systems are designed around tapes and having to rotate tapes. Some even continue to use virtual tapes to facilitate backing up to disk. If you're backing up to disk, as opposed to tapes, you need to bend your head around how to manage these tapes and this rotation schedule. It can be complex and bulky to manage.

Veeam and R1Soft CDP both address that and have a concept called a recovery point. So you have your full initial backup, and all the subsequent backups are just block-level deltas of what is to be retained. The recovery point concept is much more efficient from a space utilization perspective. And it's also much easier to manage without really complex spreadsheets and logic flowcharts. You can just say, "Hey, I want 30 daily backups of this server, six monthly backups, and three yearly backups." It's very simple and straightforward.

Besides older backup technology, the other problem we run into with backup systems is with file-based systems, as opposed to block-level systems. Why is that? On servers that have millions of files, we find that the performance of the file-based backup systems goes way down.

Bacula is an interesting example of this. We did a proof of concept and found that for our workloads, Bacula just didn't work well for most of our customers. This is because typical web hosting servers have millions of files.

Every time a web hosting server like that runs a backup, it has to go through and enumerate these files to see if there's any change since the last time the backup happened. There's just a huge amount of overhead with file-based backup systems.

Biggest Mistake IT Departments Make with Data Backups

What is the biggest mistake internal IT departments make with data backups? People assume that backups are "set it and forget it." The reality is you constantly need to be monitoring your backup servers and policies to make sure they're running regularly and reliably. Why? Things can and do change.

We have an entire team whose primary job is to manage the backups on our network. If you're on a small IT team, or even a big IT team that's resource constrained, it's easy to say "I'm not going to worry about the backup system today. I think it's running and it's okay. I've got bigger fish to fry."

Along the same lines, another huge mistake: no one ever tries to restore the data and test the restore. From this process, you should be able to answer:

- How long is it going to take to restore this data if we need it?
- Is the data going to restore in a way that's usable to us?
- Or are we going to have to jump through hoops to make it available again?

If you don't know the answer to these questions prior to needing to restore from backups, you're probably in for a long night or even a long week.

Estimating Data Backup Costs

What should data backup cost and what do those cost estimates depend on? We have a lot of ways to deploy both R1Soft and Veeam in cost effective scenarios.

At the low end is our R1Soft CDP package that includes 100GB of backup space for one system. This package is $120 per year. Veeam is much more expensive due to the Veeam licensing fee. The same amount of space on Veeam costs $240 per year.

At the high-end, we have customers who have dedicated backup infrastructure. These customers pay us a monthly fee based on their needs.

Let’s say a customer has a server with 8TB of storage. That’s one fixed fee.   Then the customer pays a per agent software licensing fee -- Veeam or R1Soft CDP for the licensing.   This is another way we provide the full solution for customers who really need to scale their backup infrastructure.

We have some customers protecting as little as 100GB of storage, and we have other customers that protect over 100TB of space. We can work with your organization to make a custom data protection solution that meets your needs and budget.

Planning the Perfect Data Protection System

Is there such a thing as a perfect data protection system? And if so, what would it look like? From our standpoint, the perfect data protection system doesn’t exist. Why?

The two biggest constraints with backing up data are speed to backup the data and the speed to restore the data. You can take all kinds of steps to optimize the speed at either the application level, the kernel level, or the network level.

But fundamentally, if you need to take 20 TB of data that’s been backed up and restore all of that, it’s going to take a long time.

On our private cloud platform, using Nimble Storage, we can mitigate these issues by using snapshots, in addition to traditional data backup tools. We can quickly and effortlessly expose snapshots to restore huge amounts of data in very little time.

You might be wondering, “why bother with backups at all if you have snapshots on your storage attached network (SAN)?” Backups are used to protect against some sort of catastrophic SAN failure event, which--though probably rare-- can and do happen.

For example, just talk to the folks that run the Australian Tax Office, which has faced a number of SAN related issues, causing major disruptions to their end users.

Data Backups vs. Disaster Recovery (DR)

What is the difference between data backups and disaster recovery (DR)? Having backups is just one part of an overall disaster recovery strategy.

Disaster recovery is about having a clearly defined way of reviving your critical IT infrastructure in case of a disaster. Backups alone don't do that. At a very basic level, when your data is backed up, sure it's protected, but nothing else is in place to make it readily available

A fully based disaster recovery solution will include most of the following elements:

- ✓ **Ability to handle failover traffic and adequate bandwidth to support peak demand**
- ✓ **Detailed Disaster Recovery/Failover Run Book with technical processes/procedures**
- ✓ **Dedicated or shared compute capacity at a DR site**
- ✓ **Development of a networking strategy to minimize reconfiguration of servers and clients**
- ✓ **Offsite backups**
- ✓ **Onsite staff to handle failover operations, including key individual who can declare disaster**
- ✓ **SAN snapshots and SAN-based replication to an offsite data center (or a storage agnostic replication technology)**
- ✓ **Well defined RTOs and RPOs**

As you can see, data backups are just one piece of a disaster recovery solution.

Disaster Recovery as a Service (DRaaS)

While we've touched on the need to couple your data backup program with disaster recovery planning, many companies need external help in the form of Disaster Recovery as a Service (DRaaS).

DRaaS is a fully-baked solution that we manage end-to-end including

- ✓ **the primary environment**
- ✓ **the failover environment**
- ✓ **the firewalls of both sites**
- ✓ **the storage of both sites and the replication**

We jointly work with the client to establish criteria, such as:

- ✓ **What exactly is a disaster that would lead us to failover?**
- ✓ **What does the process look like when we do a planned or an unplanned failover?**
- ✓ **What does the process look like when we do a planned or an unplanned failback?**
- ✓ **What does the process look like when we do all of these things?**

Every mature organization has disaster recovery requirements. A lot of people can be intimidated by the complexities of just managing infrastructure that spans multiple states and needing to constantly monitor things like SAN replication and snapshots.

Sometimes companies that are offering DRaaS are using third-party vendors. That's not always a bad thing. And for some companies, it can be a very viable strategy to have someone that manages your DR environment in Microsoft Azure for example.

However here's the key point: You don't have disaster recovery unless you routinely test your disaster recovery plan. In major enterprises, that testing can mean teams of 20 or 30 people locked in a conference room for a weekend, walking through the steps, and validating something every step of the way. For smaller sized organizations, it might be a much simpler process, much less time consuming, and involve far fewer people.

Make sure that it's possible to get the DR site up and running, and functional in the (Recovery Time Objective) RTOs and (Recovery Point Objectives) RPOs that have been defined in your company's business requirements.

Testing Disaster Recovery

How often should your company test their recovery solutions? On an annual basis? Quarterly basis?

We suggest to our clients that we do a full failover event at least once a year, assuming their operations can allow such disruptions.   We also recommend doing a soft failover once a year as well. During this kind of testing, instead of actually failing over all the workloads for the DR site, we:

- Make the storage snapshots of the DR site available
- Bring everything up
- Make sure it's working the way you would expect it to be
- Shut it back down
- Restore the replication

This soft failover approach allows us to validate that the DR plan generally is working as it should be, but it often doesn't bring to light odd corner case scenarios that a full failover test will.

Without having the capacity building experience of performing soft and hard failover tests, the key people involved in the process will lack the muscle memory of having done a failover in a controlled, minimally stressful environment. If and when a real disaster strikes, key staff members may make critical mistakes that delay the failover process.

Getting Buy-In from Your Board of Directors

What's the biggest benefit of DRaaS to CIOs and internal IT departments? The biggest benefit of investing in a solid DRaaS program is having a clearly defined plan in the event that disaster comes, so everybody knows what to do.

This way your company can put critical people back to work, bring production systems back online, serve customers, rescue any lost revenue stream that would have otherwise been halted, and avoid long-term financial harm to your company.

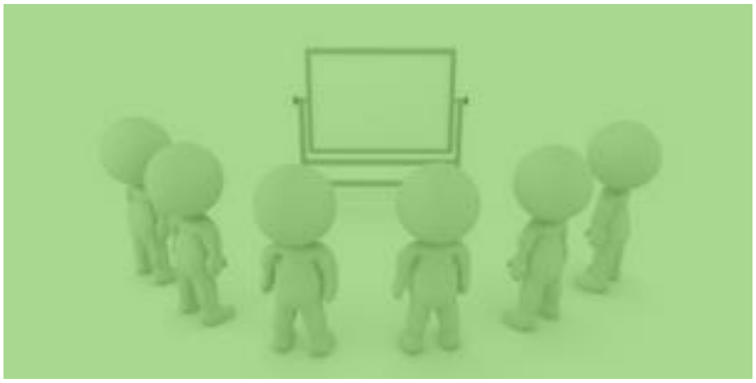

This all seems like reason enough for every company to invest in some kind of DRaaS program. But what if your company's leadership team, investors, or board of directors doesn't see value in disaster recovery?

There seems to be a lot of parallels with high-profile security breaches, where many boards lack the cyber security expertise to know what's going on, understand the value of protecting IT infrastructure, and grasp what's at stake.

We believe that there's simply a lack of education and understanding. The CIO is trying to communicate with the investors and board members, but it may not be getting through.

Given the sequence of events we've had recently with cyber-security breaches having far-reaching implications, those are great use-case scenarios to say, "look, this does happen." Otherwise, people sit back in denial and say, "well, that won't happen to us." And then when it does, ultimately, it's too late.

Estimating Disaster Recovery Cost

What should DRaaS cost? And what does it depend on?

Budgeting for DRaaS depends on so many different factors. It's hard to pin it down. There are a number of factors to consider:

- How quickly are you looking to replicate?
- What is your Recovery Point Objective (RPO)?
- What is your Recovery Time Objective (RTO)?
- Are your RPOs and RTOs different for certain applications?
- What quantity of data is associated with your RPOs and RTOs?

Next Steps

CIOs and IT directors have a lot on their plates. When teams are understaffed, data backup and disaster recovery are often early casualties.

In this eBook, you've been introduced to popular data backup platforms, vertical market considerations, common problems people make with data backup, and estimating data backup costs.

However no matter how thorough your data backup program, that alone will not help your company recover from potentially catastrophic digital and business meltdowns – which is where Disaster Recovery as a Service (DRaaS) comes into play.

In the context of disaster recovery, you've learned what DR is, how to test it, how to avoid common DR problems, how to get buy-in from other stakeholders, and how to begin estimating costs.

HOW DOES YOUR DATA BACKUP AND DISASTER RECOVERY STACK UP?

Learn more about how your current data backup and disaster recovery stacks up against industry best practices.

Do you know where some of your biggest, most dangerous gaps exist?

Sign-up for a complimentary Data Backup and Disaster Recovery Consultation.

http://content.handynetworks.com/dbdr

How Do You

COMPARE?

ABOUT HANDY NETWORKS

Backup and Disaster Recovery Planning Guide for CIOs

Written by
Jay Sudowski, Co-Founder & CEO at Handy Networks

With contributions from
Jeff Shotnik, Systems Engineer at Handy Networks
Matt Sudowski, Key Accounts Manager at Handy Networks

Handy Networks is not a typical data center company. Established in 2000 and operating under a previous trade name since 1997, we are stable, trusted hosting company that provides a variety of colocation, dedicated server hosting, managed server hosting, managed security services, and cloud hosting services.

Our clientele is largely international and we have customers from over 40 countries in the world. We have a small, but talented team of IT professionals and technology visionaries that work hard every day to ensure we meet our mission to our clients, which simply put is the following:

- **Quality** – We deliver a quality service at a competitive price.
- **Value** – We value and take pride in the relationships that we have with our clients.
- **Partnership** – We develop true partnerships with our clients. Your success guarantees our success.
- **Transparency** – We are open and honest with our customers, always.

Your Managed Cloud Hosting Provider That Cares

Backup and Disaster Recovery Planning Guide for CIOs

HANDY Networks LLC