Veeam® Software Special Edition

DRaaS

FOR DUMMIES®

A Wiley Brand

Brought to you by

veeam

Peter Gregory

About Veeam

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Enterprise™, a business that must operate 24.7.365. Veeam has pioneered a new market of Availability for the Always-On Enterprise™ by helping organizations meet recovery time and point objectives (RTPO™) of less than 15 minutes for all applications and data, through a new solution that delivers high-speed recovery, data loss avoidance, verified recoverability, leveraged data and complete visibility. Veeam Availability Suite™, which includes Veeam Backup & Replication™, leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs, while always supporting the current and future business goals of Veeam customers.

Founded in 2006, Veeam currently has 45,000 ProPartners and more than 230,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. Visit **https://www.veeam.com.**

DRaaS

FOR DUMMIES

A Wiley Brand

Veeam® Software Special Edition

DRaaS

FOR

DUMMIES®

A Wiley Brand

Veeam® Software Special Edition

by Peter Gregory

FOR

DUMMIES®

A Wiley Brand

Publisher's Acknowledgments

Table of Contents

Introduction

· ··

Disaster recovery (DR) planning has a reputation for being difficult and time consuming. Setting up alternate processing sites, procuring hardware, establishing data replication, and failover testing have been incredibly expensive undertakings. To top it all off, the need for 24x7x365 business application availability threatens to make disaster recovery planning an exercise in futility.

Disaster Recovery as a Service, or DRaaS, is turning the DR business on its head. The responsibility for all of the gritty details one used to have to juggle in order to ensure that every system, file, database record, and network element was duplicated at an alternate processing site can now be passed onto a trusted service provider. A face — not just an interface.

DRaaS is bringing true DR capabilities to an entirely new pool of organizations — folks who previously considered DR to be out of their reach. Today, DRaaS makes setting up DR almost as easy as setting up a new smart phone — seriously. People who set up DR using DRaaS are amazed at how much knowledge they *aren't* required to have. This makes DRaaS available and attractive to an even larger audience.

DRaaS For Dummies lays the foundation for this new approach to DR. After reading this book, you'll have a new appreciation for DR professionals and how difficult it used to be for them.

About This Book

Regardless of your role in the organization, taking steps to ensure that critical IT systems will survive a disaster matters. Your role determines the part you play when it comes to extending availability of critical applications to the cloud through DRaaS.

If you are a CIO or IT Director, you already know the importance of disaster recovery planning and how complicated it can be. A disaster can strike at any time, and survival of the business may well depend on the quality and accuracy of the DR procedures that took so much effort to produce (and more to maintain). Know that there is an easier way to do the technology portion of DR through DRaaS capabilities.

If you are a CISO or Security Manager, this book gives you added confidence that your IT organization may at least be willing to undertake what might be your company's first truly successful DR effort. DRaaS makes it all possible.

If you are in general business management, with this book under your belt you'll be in a better position to understand how disaster recovery planning actually works. It should be apparent to you that DRaaS solutions make part of disaster recovery planning and testing considerably easier than it used to be.

How This Book Is Organized

The primary purpose of this book is to discuss recent innovation in disaster recovery services, and how these services can be used as the core of a disaster recovery plan that is far easier to set up and test than plans using older, manually operated technologies.

Chapter 1: Understanding Disaster Recovery as a Service

Chapter 1 reviews today's disaster recovery practices as well as the changes that DRaaS is bringing to organizations. This is followed by discussions of backup and replication, the ways in which virtualization technology facilitates DRaaS, and the benefits from using DRaaS.

Chapter 2: Defining Your Requirements for a DRaaS Solution

Chapter 2 continues the discussion by explaining in detail how a disaster recovery plan is developed. The process begins with a Business Impact Assessment, followed by the establishment of key recovery objectives. This is followed by a discussion of DRaaS in public cloud, private cloud, and service provider settings. The chapter concludes with a (brief) discussion of the legal and compliance issues associated with DRaaS.

Chapter 3: Implementing DRaaS

Chapter 3 lays out the steps needed to implement DRaaS, starting with the selection of a service provider, developing a DR plan, and configuring resources. For enterprises with their own data centers, there is a discussion on the use of DRaaS tools internally.

Chapter 4: DRaaS Operations

Chapter 4 discusses operational aspects of a DRaaS system, including network monitoring, conducting failovers and failbacks, and testing a DR plan. The chapter closes with a short discussion on switching service providers.

Chapter 5: Ten Business Benefits of DRaaS

Chapter 5 contains the ten most important business benefits that organizations using DRaaS will enjoy. At least some of these benefits — and perhaps more than just a few — surely await you in your DRaaS future.

Icons Used in this Book

Icons are used throughout this book to call attention to material worth noting in a special way. Here is a list of the icons along with a description of what each means:

Some points bear repeating, and others bear remembering. When you see this icon, take special note of what you're about to read.

Watch Out! This information tells you to steer clear of things that may leave you vulnerable, cost you big bucks, suck your time, or be bad practices.

This icon indicates technical information that is probably most interesting to technology planners and architects.

If you see a Tip icon, pay attention — you're about to find out how to save some aggravation and time.

Where To Go From Here

Regardless of whether you are still considering developing a DR plan, or you're already relying on DRaaS for your main DR plan, this book will help you understand more about the benefits of DRaaS, how easy it is to set up, and the benefits your organization will enjoy as you protect your organization's critical data and services from disasters.

Chapter 1

Understanding Disaster Recovery as a Service (DRaaS)

··

In This Chapter

▶ Understanding today's disaster recovery needs

▶ Introducing disaster recovery as a service (DRaaS)

▶ DRaaS and virtualization

▶ The benefits of DRaaS

··

Disaster recovery (DR, for short) is the undertaking whereby an organization invests in computing hardware and software to be used in the event that a disaster renders the primary processing site unavailable. That's about as simply as I can describe it, but in reality it is far more complex than that.

People in today's always-on, always connected world are far less forgiving of unscheduled downtime that occurs, regardless of the reason. In this world, application availability is king. Not that long ago, people tolerated applications being down for hours or even days at a time (considering the circumstances, of course), but today even a fraction of an hour is considered inexcusable. We want our application and we want it now!

In this chapter, you'll get a chance to take a look at application availability expectations, the mechanisms that comprise

DRaaS, and the benefits that organizations can enjoy with DRaaS solutions.

Today's Disaster Recovery Practices

It's not easy being a CIO today. CIOs are under pressure to make their applications and data available continuously, without regard for the "stuff" that happens: Hardware failures, software bugs, data corruption, and disasters. In today's point-and-click world, business users think it's easy for an IT organization to create a fault-tolerant, disaster-proof environment. CIOs and others in IT know it's anything but.

Enter disaster recovery — DR for those of you who love acronyms. Traditional approaches to DR include hot site, cold site, and warm site, discussed here.

Hot site

In a hot site approach, the organization duplicates its entire environment as the basis of its DR strategy — an approach which, as you'd expect, costs a lot in terms of investment and upkeep. Even with data duplication, keeping hot site servers and other components in sync is time consuming.

A typical hot site consists of servers, storage systems, and network infrastructure that together comprise a logical duplication of the main processing site. Servers and other components are maintained and kept at the same release and patch level as their primary counterparts. Data at the primary site is usually replicated over a WAN link to the hot site. Failover may be automatic or manual, depending on business requirements and available resources.

Organizations can run their sites in "active-active" or "active-passive" mode. In active-active mode, applications at primary and recovery sites are live all the time, and data is replicated bi-directionally so that all databases are in sync. In active-passive mode, one site acts as primary, and data is replicated to the passive standby sites.

Cold site

Effectively a non-plan, the cold site approach proposes that, after a disaster occurs, the organization sends backup media to an empty facility, in hopes that the new computers they purchase arrive in time and can support their applications and data. This is a desperate effort guaranteed to take days if not weeks.

I don't want to give you the impression that cold sites are bad for this reason. Based on an organization's recoverability needs, some applications may appropriately be recovered to cold sites.

Another reason that organizations opt for cold sites is that they are effectively betting that a disaster is not going to occur, and thus investment is unnecessary. I don't think this is a smart move.

Warm site

With a warm site approach, the organization essentially takes the middle road between the expensive hot site and the empty cold site. Perhaps there are servers in the warm site, but they might not be current. It takes a lot longer (typically a few days or more) to recover an application to a warm site than a hot site, but it's also a lot less expensive.

Comparing hot, warm, and cold

The trouble with all of these hot-warm-cold approaches is that they do not meet today's demands for cost effective and agile recovery. Users typically expect applications to be running within a fraction of an hour. Engineered correctly, a hot site can meet this demand, but at spectacular cost. Warm and cold sites don't even come close.

It should not come as a surprise to you that most organizations "go commando" with regards to their DR plans. They have little or nothing in the way of policies, procedures, or technologies that enable the recovery of critical systems at any speed. This is understandable, as rapid recovery capabilities have historically been so expensive that only the largest organizations could afford them.

Backing up your data

Data backup is an essential part of sound IT management. We all know that things occasionally go wrong in IT, and data loss is a result that no one will tolerate.

Better organizations employ the 3-2-1 rule when it comes to backing up data. Here is how the rule works:

- ✔ Keep 3 copies of data: 1 primary, 2 backups
- ✔ Use 2 different types of media
- ✔ Keep 1 set in the cloud in DRaaS or BaaS (backup as a service)

Introducing Disaster Recovery As A Service

Since the early 2000's, many types of service providers have emerged and built entire industries that reduce the cost and complexity of many classes of technology. For instance, Software as a Service (SaaS), Infrastructure as a Service (IaaS), and Platform as a Service (PaaS) have created entirely new paradigms for businesses' use of technology.

Disaster Recovery as a Service (DRaaS) is a rapidly growing cloud-based service that makes it easy for organizations to set up alternate processing sites for disaster recovery purposes. Like other "as a service" offerings, advanced software enables DRaaS to simplify the entire process for organizations of any size as well as the service providers that offer this service.

DRaaS is important because it represents an innovative and less costly way to back up critical data and quickly recover critical systems after a disaster. DRaaS does this by leveraging cloud-based resources that provide infrastructure that is far less expensive than on-premise systems due to the ability to scale and share cloud resources

To meet the growing demand for software resilience, DRaaS has brought simplification and reduced costs to organizations that are serious about implementing DR. With DRaaS, an organization can implement a high-performing DR solution for its

critical systems but without any of the complexities. Like other "as a service" providers, DRaaS providers take care of the back-end complexity for their customers and provide a simple user interface for setting up and managing a DR solution.

There are a few different flavors of DRaaS discussed here, related to whether your organization uses public or private cloud.

Public cloud DRaaS

Organizations can implement DRaaS using a public cloud infrastructure. Any public cloud service that meets an organization's security and operational requirements can be used. A typical DRaaS solution will employ customer-managed software for setting up and controlling cloud based DR resources.

While their vast scalability provides many cost advantages, going with a public cloud infrastructure means you will likely be foregoing a personal one-on-one relationship. If something goes seriously wrong — get ready to stand in line, if you can find a line to stand in!

Private cloud DRaaS

Organizations with their own data centers and private cloud infrastructure can definitely utilize DRaaS solutions. The software components that comprise DRaaS solutions can be installed on an organization's own server infrastructure. In these types of situations, the HQ datacenter in essence takes on the role of the service provider for their different business locations. These solutions will reduce the effort and complexity of data backup and replication mechanisms for organizations that are required to keep data under their direct physical control.

Managed cloud DRaaS

Organizations using managed cloud services can include DRaaS solutions to their service portfolio. Managed cloud service providers can include DRaaS as a part of a standard, hands-free offering that takes care of data backup and data replication details. This permits customers to concentrate on their software applications and other hosted components.

The Role and Need for Secondary Sites

One of the time honored (and still valid) principles of disaster recovery planning states that a secondary computing location be established. The reasons for this include

- ✔ The primary site may be incapacitated because of the effects of a regional disaster. This includes events such as an earthquake, hurricane, or flood.

- ✔ The primary site may be incapacitated by the effects of a localized event, such as a fire, landslide, power failure, communications outage, or a water main break.

- ✔ The primary site may have suffered an equipment failure in its IT infrastructure, or an operational error resulting in unexpected and perhaps prolonged downtime.

The best bet for covering all of these scenarios is the use of an alternate processing center some distance away from the primary site, generally 100 miles or greater, depending on the types of disasters that can happen in your part of the world. This helps to ensure that the alternate processing site is not affected by whatever regional event has affected the primary site. This approach is still valid with cloud services. With a cloud-hosting provider, you'll generally have a choice on where your recovery servers will reside. What you don't want to end up with is a situation where the DR servers assigned to you are in the same city as your primary site. This would not result in a good recovery scenario, since the hosting provider may be adversely affected by the same disaster that affects your primary site.

Using a cloud-based hosting provider is a cost-effective way to build a secondary site. The main advantage is the preservation of capital. Virtually no investment in recovery systems is required, since they are instead leased from the service provider if and when they are needed.

Backup and Replication

An essential part of a disaster recovery plan is some means for transporting copies of mission critical data away from the primary processing site to another location that will not be affected by whatever event affected the primary site. There are two main ways to copy data:

✔ **Backup.** Data is copied from databases, flat files, and virtual machine images to backup media residing on disc-based storage, but could also include backup to magnetic tape or virtual tape libraries.

✔ **Replication.** As data is being written to databases and flat files, that same data is being transmitted over a network to another storage system, usually to an alternate processing center or cloud provider.

The main distinction between backup and replication is this: Backup copies the entirety of a machine image, files, or databases (or the incremental changes since the last backup), in a one-time operation that is then repeated periodically; whereas replication is the continuous or near-continuous transference of updated disk blocks — say, batch updates every five minutes.

Backup was once considered "good enough" for disaster recovery purposes. However, *good enough* implied that an organization was willing to wait days to recover their systems and get them running again. However, in today's always-on enterprises, backup is no longer good enough: Backup and replication together are necessary for organizations of all sizes to get its critical applications up and running in 15 minutes or less.

The right strategy for today's DR needs, then, requires both backup and replication: Frequent backup of virtual machine images and the replication of critical data. Together, these provide system recovery synergy that facilitates rapid restoration of critical systems.

DRaaS and Virtualization

Virtualization — the technology that permits multiple operating system instances to run on each physical server — has freed up IT infrastructure and facilitated the revolution that is the mass migration of applications to the cloud. Individual operating systems (which reside within *virtual machines*) reside in "images", which are large flat files that can be copied to a DR site for rapid recovery of servers. The power of virtualization and virtual machine management have contributed significantly to the power of DRaaS.

Some DRaaS solutions have the ability to provide advanced, imaged-based virtual machine (VM) replication, which can be used to send VM images to a cloud service provider. Service providers can provide virtual cloud hosts; recovering your server is as easy as booting those hosts from the images sent from your primary site.

Better DRaaS solutions include *agentless* components — meaning there is no software present within individual virtual machines. Instead, you install a module in the virtual environment, which intercepts local disk traffic and sends it to your cloud service provider, where another module receives the traffic and keeps your server VM's and databases up to date, usually within minutes.

DRaaS Benefits

DRaaS represents the next generation of rapid system data recovery and always-on availability, helping organizations avoid downtime and business disruption without the high costs associated with traditional hot sites.

The low cost and simplicity of DRaaS makes it available to an entirely new class of organizations. The ability to recover applications in the cloud, if and when needed, slashes the cost and complexity of recovery capabilities. Organizations that were on the sidelines, longing for DR capabilities, can now enjoy capabilities that were once reserved for large organizations.

Chapter 2

Defining Your Requirements for a DRaaS Solution

· ··

In This Chapter

▶ Using a Business Impact Assessment to identify critical processes and servers

▶ Setting recovery targets

▶ Working with public cloud, private cloud, and local service providers

▶ Working through the legal and compliance issues associated with DRaaS

▶ Understanding how DRaaS untangles network complexities

· ··

*P*lanning a DRaaS solution is one of the easier projects that an organization will undertake. More than likely, the most difficult parts won't have anything to do with the DRaaS technology per se, but rather will come from having to first decide which applications and databases will be in scope, and then tracking down where all of their data resides. In this chapter, I take you through the concepts and steps needed to get your DR plan rolling.

Performing a Business Impact Assessment

Before you jump into your catalog of enterprise applications and look at how large their databases are, it's best if you step back and assess which business processes and applications

are truly the most important in your organization. Enter the Business Impact Assessment, or BIA. I'll give you the short-hand version for a BIA here; for a more complete recipe, see my book, *IT Disaster Recovery Planning For Dummies*.

The objective of a BIA is to identify those business processes — and their underlying IT systems — that are truly the most critical for the organization. You start by listing them, in a spreadsheet preferably, row by row. In columns to the right, list the IT systems and data stores each process is dependent upon. In more columns, list dependencies: The processes and systems that each process is dependent upon, and the processes and systems that depend on this one. Most find a tangled web of dependencies. In that cloud is a silver lining: With some analysis you'll find opportunities to simplify and correct dependencies that ought not exist.

Next, you make a guess (or seek input of key IT and business stakeholders) on the total downtime that the organization can tolerate for each business process before the business itself fails.

You might think of other columns to add, and then you can sort and filter the list any way that you like so that you can get to the best reckoning of those processes and systems that are the most important in your organization. Congratulations! You have just created the scope for your DRaaS project. (Well, maybe. I did say this was an easy recipe, like bake-less brownies, that are okay for around the house, but not fit for a banquet.)

Setting Recovery Targets

When you have completed the Business Impact Assessment and established your list of critical processes and systems, the next task before you involves setting recovery targets. These targets determine how quickly you want to have your site running after a disaster and how much data loss you can live with.

The two main recovery targets are

✔ **Recovery Time Objective (RTO).** This is a Service Level Agreement (SLA) setting the maximum period of time that an application will be unavailable in a downtime

situation. The shorter the period of time, the faster the recoveries will need to be. Generally, this data includes everything about the operating system, plus applications, configuration data, as well as application data.

✓ **Recovery Point Objective (RPO).** This is an SLA setting the maximum amount of data loss that would occur in a downtime or data loss scenario. Frequent backups or near-continuous replication is required to reduce the amount of data loss.

Recovery time and recovery point objectives are related, although they do not have to have the same values. Here are a few examples:

✓ **RTO = 15 min, RPO = 15 min.** System will be back up at DR site in < 15 minutes, with no more than 15 minutes of data loss. Application data and OS data (or VM images) are copied more frequently than every 15 minutes, or replicated.

✓ **RTO = 60 min, RPO = 10 min.** System will be back up in < 60 minutes. The maximum data loss is 10 minutes.

Supporting Small Businesses with DRaaS

Smaller organizations have less data and less infrastructure to contend with. This makes their DRaaS environments simpler and easier to manage. But smaller organizations also have fewer IT specialists, and few (if any) experts on any given IT topic. This, however, makes better DRaaS solutions a real boon to smaller organizations: because they are easy to plan, build, and manage, that deep expertise isn't needed.

Smaller organizations will probably want to utilize the services of a managed cloud provider (one with a telephone number and actual people to work with), versus the larger public cloud providers that assume that customers have all the necessary expertise.

Supporting Large Businesses with DRaaS

Larger organizations have more data and — unsurprisingly — more infrastructure. Better DRaaS tools that show all of the data- and DR-related resources in an easy-to-use interface simplify the planning and management of DRaaS capabilities. Larger organizations feel the resource squeeze; better DRaaS solutions simplify all aspects of DR capabilities, enabling the smart people to focus on other things.

Cloudy with a Chance of DRaaS

There are a lot of options available to you when planning your DRaaS solution. Whether you want to use public cloud, private cloud, or partner with a managed cloud provider, DRaaS can be set up to back up and replicate your critical data and system images to facilitate an easy and rapid recovery.

DRaaS with public cloud

When I talk about public cloud providers, I mean any of those market leaders in the cloud computing field — Amazon Web Services, Microsoft Azure, and many other well-known brands. These are the mega-players with nearly limitless resources and scads of locations around the world.

DRaaS takes a little more effort to set up with public cloud providers. The reason for this is that you'll be spending more effort setting up your DRaaS solution since you will have to architect and configure the cloud side of the system in addition to settings within your own environment. However, if you are already committed to doing your DR with a big public cloud provider, DRaaS can definitely work for you.

DRaaS with a local service provider

Think of a local service provider as a "co-lo" or hosting provider with the genuine customer service feel (as in, they

have a phone number and you can talk with a human) versus public cloud companies that lack the personal touch. Like we have discussed before, the face, not just the interface.

Hands-down, partnering with a local service provider is going to be your best bet for a DR solution that you can depend on during a disaster situation.

While you can replicate and fail over to a public cloud, it's a little more complicated to set up and maintain, more like a DIY project where you assemble all of the pieces yourself (think Scandinavian furniture). The advantage with local service providers is that they often provide the interfaces for you that you just plug into. This can be really helpful when you are setting everything up, as well as during troubleshooting.

If and when you have a disaster, stress levels will be high and you will likely need additional help; then is when a local service provider will shine and be there for you when you need it the most. A local service provider can also provide expertise you may not have in your organization, including big picture DR planning and DR best practices.

DRaaS with private cloud

If your organization already has multiple processing centers in diverse locations (which you may refer to as a type of private cloud), you can still utilize DRaaS solutions to simplify your DR capabilities. With private cloud, you might not be enjoying quite the economy of scale that you would with a public cloud provider, but then again maybe you opted for private cloud because you still need some measure of control. No matter. DRaaS is possible with this type of configuration, with someone internally taking on the role of the service provider for the diverse locations.

Identifying Security and Compliance Issues

Self-driving cars relieve us of a lot of details, but we are still in charge of where we're going. With DRaaS, you're simplifying the process by relying on the expertise of others, but you are

still in the driver's seat — and are still responsible for what happens on the road. Whenever you send your data out of your physical control, there are several security, regulatory, and other legal issues you may need to be aware of.

Data protection

A lot of laws and regulations require that organizations enact several forms of security controls to protect sensitive data. Some of the measures required include encryption of stored data, logging and monitoring, and strong access controls. These requirements apply, regardless of how or where organizations process and store sensitive data.

Some cloud providers include some of these security capabilities as additional features, and certainly they may be a key differentiator that will help an organization choose which service provider to use. Just be sure that these features work as intended and don't get in the way of the ability to operate critical applications in the cloud.

Geography and the law

With regards to sending your critical data to a service provider, you have to know — with pretty good precision — where the data is actually going, and where else it might go in different circumstances.

The reason for this is related to a legal topic called data sovereignty. What this means is that the legal owner of data, as well as the country whose citizens' data is involved, must have a measure of visibility and control over personally identifiable information. Here are a few examples:

- ✔ **U.S. state Personally Identifiable Information protection laws.** Many of the various U.S. state laws on the protection of Personally Identifiable Information (PII, for short) place certain obligations on the organizations that store and process this information. Often the jurisdiction defined by these laws extends to the organization that sent the data, even if they are not in the same state.

- ✔ **European Union privacy rules.** Several European countries place conditions and even restrictions on the

cross-border movement of PII, especially if the data is leaving the EU altogether.

✔ **U.S. PATRIOT Act.** Many European countries, as well as Canada, have grave concerns over the provisions of the U.S. PATRIOT Act, particularly with regards to federal agencies that can seize data and forbid the data custodian from saying anything about it.

REMEMBER

Some governments may claim that their privacy laws apply to you, even if you don't have business locations in those states or countries.

Additional legal and regulatory Issues

Geography aside, there may well be other issues regarding the transmission of data to another organization. There may be contractual obligations between an organization that wants to do DRaaS and the organization's customers, which place obligations and restrictions on what the organization can do with its customers' data.

Legal issues you'll need to consider may include the following:

✔ **Data ownership.** From a regulatory as well as a contractual perspective, it is important to know who is the legal owner of the data you store and process.

✔ **Data custodian.** Essentially, this means who is the responsible party for protecting this data from compromise, unauthorized access, unauthorized alteration, and loss.

✔ **Data access.** You need to understand any legally imposed rules concerning access: Who is permitted to have access to the data, in other words. This may include parties that have physical access to storage systems containing data, like service providers.

✔ **Data use.** While not directly related to DR, it's important to know what business uses are permitted for the data you store and manage. For instance, you may be processing patient information, but you might not be permitted to market to those persons, even if you have the capability to do so.

Insurance

Insurance plays a key role in IT operations. Actuarial data for breaches and other events are improving, and more insurance companies are getting into the game with more offerings.

Insurance companies are accumulating data on the factors that lead to various types of loss events. Companies that meet certain requirements are enjoying lower premiums and better loss payouts. In many cases, the enactment of controls and processes results in lower premiums. These factors vary from insurance company to insurance company, but over time the common themes will emerge. The degree of preparation an organization has for a disaster may (if not already) be one of those factors recognized by the insurance industry.

✔ **Data retention.** Organizations are often required to retain certain sets of data for arbitrary minimum or maximum periods of time. For instance, a financial service provider may be required to retain its financial records for seven years but no more than ten years.

✔ **Data location.** You may be obligated to store the data within certain jurisdictions or geographic areas or, conversely, prohibited from storing or processing the data in certain locations. When storing data with a third party such as a cloud provider or local service provider, you will need to understand how the provider stores data in its various locations, and what visibility or control you have over this.

Your legal obligations may include the need for periodic internal or external audits, so that you can demonstrate your compliance to your obligations. Often you will be asked to demonstrate your compliance to any or all of your requirements.

Untangling Network Complexities

Networks within your data centers and between your data centers are an important part of your DR capabilities. In this section I discuss networking issues that you can't overlook as you develop your DR plans.

Duplicating network architecture

Early DRaaS solutions were relatively simple for small organizations with a simple, straightforward network infrastructure, because recreating a simple network design in the cloud was fairly easy.

Today's DRaaS capabilities have advanced to the point where even a complex network can be replicated in the cloud. Instead of just failing over a few critical servers to the cloud, all of the servers, together with their supporting complex network infrastructure, can be run in the cloud in a disaster scenario. This makes DRaaS more feasible for larger organizations with their complex networks.

Network performance

One of the most critical factors related to the architecture of your DRaaS solution is the performance of the network(s) between your primary site and your DR site(s). The two main ways to measure network performance are

✔ **Latency.** This is a measure of the time it takes for a single packet to travel from your primary site to your DR site. Generally, the greater the distance that your data has to travel, the greater the latency will be. This is part of the reason why it is important to select a DR site that is far enough away to avoid involvement with whatever regional disaster affects your primary site, but not too far away as to cause excessive latency.

✔ **Bandwidth.** This is a measure of the amount of data that can be sent through the network at any given time. The volume of data you need to move from your primary site to your DR sites will determine the bandwidth required. You need to consider peak processing times and seasons and decide whether you are okay with your replication falling a little behind or not.

WAN optimization, deduplication, and acceleration tools may be able to speed up your site-to-site network traffic. These tools are not inexpensive, but they may be a better solution that using bigger or faster networks.

Bootstrapping Recovery Sites

The biggest data movement of any DR plan is the initial copy of all the involved virtual machines. To ease this process, which could be time consuming for large organizations lacking sufficient bandwidth, one could opt for a cloud seeding strategy, where the initial backup is saved to a storage device and then securely posted or delivered to the service provider.

To aid in this process, Veeam offers integrated backup and replication technology. Integrated backup can be used to create a complete copy of the customer environment, which can then be sent offline to the recovery site. The service provider can then load the copies into the DR environment, and from here the customer only needs to replicate the daily changes happening in his environment. This dramatically reduces the bandwidth consumption.

Veeam has WAN optimization built into its products, which may eliminate the need to purchase separate and potentially costly components from a cloud provider.

Chapter 3

Implementing DRaaS

· ··

In This Chapter

▶ Selecting a DRaaS service provider

▶ Building a DR plan that will ensure your organization survives a disaster

▶ Becoming an internal DRaaS provider (for larger organizations)

· ··

*A*fter you have completed the planning phase, it's time to implement your DRaaS solution. The most important aspect of this is the selection of a trusted service provider. There are a plethora of service providers to select from. The best strategy is to do your own research and narrow a list down to a manageable number, then contact more than one and interview them to see which one(s) are a best fit for your organization.

Larger organizations may opt to become their own DRaaS internal service providers. I discuss this at the end of this chapter.

Veeam has an extensive list of trusted service providers offering DRaaS, found at `www.veeam.com/service-provider-lookup.html?vcp-type=draas`.

Selecting A Service Provider

Selection of a service provider is probably the most important decision you'll make in your DRaaS project. Your service provider should play the role of your trusted partner whose deep experience with DR planning and familiarity with a deep expertise in DRaaS means you'll be sure to end up with a DR

plan that has the best possible chance of keeping your business available when a disaster occurs — big or small.

Some of the considerations you'll want to examine include these:

- ✔ **Locations.** Does the service provider operate in a single location, or does it have multiple locations available to support you? Where exactly will your data reside? Will it be within, or outside, your specific country? Often it is preferred to keep data in-country due to security and legal needs.

- ✔ **Performance.** In every way that matters, you need to know whether a service provider's infrastructure is up to snuff, but that is not the only performance benchmark. What are their RPO and RTO SLAs? How confident are they in meeting these SLAs? Have them provide real, recent examples. You also need to know how responsive the service provider itself is when it comes to questions and issues. If you call them, how quickly can you get a live voice? During what hours?

- ✔ **Capacity.** You need to evaluate each service provider's ability to transmit, receive, store, and deliver data, and determine if they have more than enough capacity to support you and all of their other clients. Can they match capacity in line with your business's data and business growth projections?

- ✔ **Shared or dedicated resources.** You need to determine how the service provider sets up resources for its customers, and the extent to which these resources are shared among them. Resource sharing is one way in which service providers make better use of those resources, but when spread too thin, no one will get what they need.

- ✔ **Testing services.** Here you need to be clear about what sorts of testing capabilities they have, including failover and recovery testing, as well as the frequency of testing

- ✔ **Experience with DR.** The service providers you consider must have specific DR experience. Otherwise, how are they going to be able to help set up a DRaaS capability to meet your needs?

✔ **Experience with DRaaS tools.** You'll want to consider a service provider that has experience with at least a couple of the leading DRaaS tools on the market. This will help them decide which DRaaS tool(s) will work best for you.

✔ **Pricing.** It's important to understand what you pay for and what you get for it. Look out for service providers that will store your data at very low costs but practically hold your data ransom if and when you need it the most by charging high prices if you actually want to recover your data.

✔ **References.** DR is too important to undertake without turning over stones to see how well a particular local service provider is doing. You'll want to speak with a few reference clients who can tell you what each service provider is doing well, and where their challenge areas are.

Working With A Local Service Provider

After you have selected a local service provider, you'll be establishing and deepening your business relationship and building trust over time. Some of the things you can expect include

✔ **Primary business and technical contacts.** These are the key people you'll be working with — your go-to people. They are there to help you directly with business and technical issues, and occasionally they will refer you to specialists. You should have not only their e-mail addresses, but their mobile phone numbers as well.

✔ **Portal.** You should be able to log into a place on the service provider's site where you can find documentation, information about the services you are using, and connections to your DRaaS resources. There should be a dashboard of some kind that will show you your costs and resource utilization.

✔ **Reports and Statements.** Your service provider should be sending you reports and statements on a regular basis. When your service provider is designing or building new capabilities for your organization, you should be getting regular project status reports as well.

✔ **User community.** You should have access to other service provider customers so that you can learn from each other, which will help everyone better utilize the service and improve their DR capabilities.

✔ **Notifications.** Your service provider should be proactively notifying you whenever they have difficulties, have made any changes that affect your DRaaS service, and whether anything has hiccupped in your DRaaS instance. Depending on the severity of the event, they should be sending you e-mails or even calling you.

✔ **Testing.** Your service provider should be on hand to assist with any and all testing that you need to do, including failover testing. It is in the service provider's best interest to ensure that your DRaaS service is working correctly for you in every respect.

Building A Successful DR Plan

Surviving a disaster is rarely an accident: Rather, a successful DR plan is required if your organization is going to be operating after a disaster. There are a lot of things that an organization needs, including

✔ **Executive support and involvement.** A DR project is not one of those activities that can take place in isolation. Instead, DR needs support from executive management, because it takes time to build a proper DR plan, and the people who best know how to do it are probably busy with other matters like day-to-day operations. Unless management mandates the DR project and its success, it's going to be difficult for participants to stay involved.

✔ **The best minds.** The people in the organization who have the deepest knowledge of company operations and IT systems are the ones who need to staff the DR team. It's no good putting the newest, greenest employees on the team; the newbies don't know enough about the details to know how critical systems can be replicated and used.

✔ **Defined scope.** In all but the smallest organizations, the scope of the project needs to clearly define which business units or departments are in scope and which are not. Otherwise, the organization is liable to bite off more

> than it can chew, putting the entire DR project at risk since it will be very difficult to complete it if too many systems are in scope. The best bet is to start small, get some wins, and expand the scope later once everyone is more familiar with how DR planning and execution work.

As I discuss in Chapter 2, it's important to start with a Business Impact Assessment (BIA) so that the required resources and dependencies for critical processes and systems can be discovered and documented. Organizations that skip the BIA may find their DR plans are incomplete because of forgotten dependencies.

Next, a very careful study of in-scope IT systems and applications needs to be undertaken. High-level concepts and details alike need to be captured, so that the organization will be able to successfully run an application at a DR site. In today's interconnected world, this usually involves strategies for dealing with the complexity of data feeds into and out of each application. These are usually at the heart of dependencies that need to be unraveled in the discovery phase.

Further, the organization must consider the very real possibility that its own core personnel may be unable to respond during a disaster. Often overlooked, sometimes in a disaster the people who best know how to run the business and its critical IT systems are unavailable because of the effects of a disaster. For this reason, it's vitally important that the organization document in rich detail every aspect of each critical business process and supporting IT system, so that other people less familiar with their inner workings can successfully function after a disaster strikes. Otherwise, all of that planning won't be worth much.

Another advantage of using a local service provider is their expertise in disaster recovery planning.

Configuring Resources

After all of the up-front planning for DR has been completed, it's time to begin setting up your DRaaS systems. By working within your on-premises data protection solution and with the tools and resources of your local service provider, configuring

your resources should be relatively quick and easy. It it's not, that might be a sign of trouble ahead.

The Veeam Best Practices Guide contains recommendations on configuring and operating your DR environment.

Becoming an Internal DRaaS Provider

Large enterprises with multiple data center locations can take on the role of a DRaaS provider, as opposed to using an external cloud provider. This is an option for organizations that have multiple data centers, and for organizations that want to maintain physical control over its data and keep everything in-house or push it out to a private cloud.

Similar to the trusted local service providers discussed earlier in this chapter, an organization doing its own DRaaS just needs to set up its backup infrastructure, establish the connections between the data centers, and then make the infrastructure available to IT personnel who manage IT systems and data.

Chapter 4

DRaaS Operations

● ●●

In This Chapter

▶ Ensuring network connections are working

▶ Executing failovers and failbacks

▶ Testing your DR plans (and not just the DRaaS part)

▶ Preparing to switch service providers

● ●●

*L*ike any critical IT system, DRaaS environments must be continually monitored and periodically tested. This ensures the DRaaS environment is working properly and that personnel understand how to use it. This chapter addresses key aspects of DRaaS operations so that you can know that it's working for you and that it will be there when it counts the most — in a disaster situation.

Monitoring Network Traffic

A healthy network connection is required for a DRaaS installation to work properly. Any organization investing in DRaaS needs to include the tools and resources to ensure that

✔ There is ample bandwidth at all times to transmit backup and replication data quickly

✔ The network is free of errors that may slow down data transmission

✔ Network traffic is free of anomalies that could be signs of errors or malfunctions

If your organization has a SIEM (security incident and event management) system, I recommend you configure all of the components of your DRaaS system so that log entries and alarms will be collected and processed.

Executing a Failover

A *failover* is an event where a workload is switched from an active machine to its replica. There are two types of failovers:

- **Planned failover:** You would execute a planned failover either as a part of a test, or because you know that one of your primary servers is about to go offline.

- **Unplanned failover:** An unplanned failover occurs automatically when a primary server has gone offline (or when a user recognizes the need and triggers a failover), then the workload is automatically switched to its replica.

Executing a Failback

A *failback* is an event where a workload is transitioned from a replica back to a production server. There are several options for a failback:

- **You can fail back to a server in the original location on the source host.** Typically, you would do this when the original server (or supporting infrastructure) has been recovered.

- **You can fail back to a server that has been restored up-front from the backup in a new location.** Typically, you would do this in a disaster scenario where the original location will be incapacitated for some time.

- **You can fail back to an entirely *new* location by transferring all VM replica files to the selected destination.** Typically, you would do this in a disaster scenario where you need to resume processing in a new location, as opposed to the original processing site.

Like a failover, a failback is a temporary state. To complete a failback, the failback needs to be committed. A failback can also be undone, returning the replica to the failover state.

As part of your DR plan, set guidelines and goals for how long you plan to run at the failover site before completing the failback. Companies often plan for the failover, but do not set goals and policies for then failing back.

REMEMBER

Since a failover is considered a temporary state, a failback is a transition to a new permanent state.

Testing Your DR Plan

Testing a DR plan is necessary to ensure that it will work in the event of an actual disaster. The very survival of your organization may be at stake, so take testing in advance very seriously.

REMEMBER

There are no second chances in disasters.

There are several types of DR testing, including:

- **Document review:** In this scenario, a number of individuals with the necessary expertise read through policy, process, and procedure documents to better understand them and to discover any errors or omissions.

- **Walkthrough:** Here, the organization conducts a guided group discussion of all DR documents, also for the purpose of discovering errors and omissions.

- **Simulation:** Similar to a walkthrough, a simulation is a guided discussion where a group leader reads through a disaster scenario that unfolds throughout the session. As the scenario unfolds, participants discuss various aspects of the response that would occur. This is designed to make the discussion more realistic.

- **Parallel test:** This is a test of IT systems' ability to handle production workloads when a failover takes place. Primary IT systems continue to handle workloads as usual; DR systems are brought online and business transactions directed to them as though they were handling production workloads. The purpose of a parallel test is to determine whether DR systems can handle production workloads and perform transactions properly.

- **Cutover test:** Here, actual production work is shifted to DR systems, to see whether they would be able to perform production work in the event of a real disaster.

Swapping Primary and DR Sites

Another approach to the cutover test is the idea of periodically swapping the role of primary and DR sites. Many organizations on the leading edge of DR execution, planning, and preparation have undertaken the task of performing a failover from its main processing site to its DR site, and then having the DR site assume full production workload for several months. The DR site becomes the new primary site, and vice versa. This is the ultimate test of confidence of a DR site, and the best way for personnel to retain their familiarity with failover procedures.

> **TIP** While cutover tests are scary for many, this is the best way to ensure you are really prepared.

Building an Exit Strategy

For any of several legitimate reasons, organizations occasionally find the need to change local service providers. For this reason, organizations adopt a mindset of keeping their options open with regards to their vendors and suppliers.

When it comes to working with — and, potentially switching — DRaaS providers, organizations will have an easier time rolling with the changes if they adopt these practices:

- **Monitor and analyze:** Organizations need to monitor their local service providers and keep track of service levels and issues. This will build a business record that will highlight each service provider's performance and help management determine which service providers are performing well and which are faltering.

- **Periodic business reviews:** Organizations should be conducting periodic business reviews with their key local service providers. This helps both parties better understand each other and ensure that they are still strategically aligned.

✔ **Maintain documentation.** It's vital that every organization develop and maintain documentation on several key facets of the business including

- Architecture diagrams

- Configuration standards

- Processes and procedures

- Data flows

Organizations with the discipline to maintain their documentation will have an easier time changing service providers, because they will have a clearer understanding of their infrastructure at a detailed level.

Chapter 5

Ten Business Benefits of DRaaS

*I*n life, as they say, the only certainties are death and taxes.

In IT, the only certainties are bugs and outages. Disasters may be a distant third, but the impact of disasters is profound: They threaten the reputation and the very survival of the organization. Without a proven DR plan, there may be little hope of surviving a disaster.

When thinking of disasters, we often think of the big ones: Katrina, Mt. St. Helens, or the 9/11 attacks. Sure, those are significant events that led to the destruction of many businesses. But those happen to other people, right? Well, yes, mostly. But there are a LOT of disasters that happen at the local level, and they take a big toll on organizations that get in the way. The kinds of events I'm talking about include

- ✔ Fires
- ✔ Water main breaks
- ✔ Hurricanes
- ✔ Floods
- ✔ Telecomm outages

REMEMBER

It's not just spectacular natural disasters that impact organizations, but the "everyday" disasters as well, when a software bug or administrator error results in data corruption or deleted objects in a SharePoint, Active Directory, or Exchange server.

These events occur on an almost daily basis someplace, and businesses are adversely affected. When an organization's main processing site or primary database is affected, this can cause hours and even days of downtime. For a smaller organization, this may be the end of the business as you know it.

Sorry about this dreadful talk. We have to face up to the bad things that can happen. But now, visualize DRaaS in your future and the business benefits that await you, starting now.

Availability

Extending your existing data center to the cloud will significantly improve the availability of your critical applications. Improving availability is the main reason organizations adopt DRaaS solutions in the first place.

Cost Reduction

Prior to DRaaS, DR capabilities were in fact a lot of disorganized pieces made to work at greater cost and effort, and required uncommon expertise to manage. Feature for feature, implementing DRaaS costs way less than traditional means, and requires no capital expenditure. There are no hardware costs — you just pay as you go.

Simplicity

A DRaaS solution brings together all of the elements of DR site planning, architecture, and capabilities into a single set of tools. Before DRaaS, doing the same thing required a lot of different tools and systems that weren't designed to work together.

This extends to your personnel: the ultra-high skills required to build and maintain more complex data replication and server failover mechanisms are not needed with today's DRaaS solutions.

Visibility

Because you'll have a single GUI for your DRaaS solution, it's going to be a lot easier for you to see what's going on in your DRaaS environment. You'll be able to see how your data backups and replicas are working, the state of all of your VM replicas, and the state of your servers (whether they are processing workloads or in failover).

Scalability

DRaaS solutions scale upward and outward to cover as many VM's, databases, storage systems, and sites that you can throw at it. With the DRaaS system, you're logically joining primary and recovery resources to each other and directing operations like replication and backup, no matter how many servers, applications, or databases are in your environment.

Like any DR capability, you've got to make sure you have enough network bandwidth for moving data, enough servers for your VM's, and enough storage for your images and data.

Flexibility

With a DRaaS solution, you're not locked into a particular server technology, database technology, backup technology, or network technology. You are free to mix and match operating systems, virtualization platforms, database management systems, and backup tools. Your DRaaS system is a lot like middleware that just makes it all work.

Easier Testing

Before DRaaS, conducting DR testing was more difficult, time consuming, and risky (in case something goes wrong,

resulting in production systems going offline). Because of DRaaS's automated management of virtual machines, backups, and replication, testing failovers is as simple as pointing and clicking.

Compliance

DRaaS helps organizations with their compliance requirements by providing controls needed to monitor and protect critical and sensitive information. You'll be able to show your auditors and regulators where your data is located and who has control over it, all through a single GUI.

Competitive Advantage

Why would your business (potentially) be better than those of your competitors? With a DRaaS solution, you'd have a far greater assurance that your business is going to be there for your customers in the future, even a future that includes a local or regional disaster.

Without DRaaS, you're implicitly betting that a disaster, even a localized one like a fire or a water main break, is not going to happen to you. Me? I prefer a safer bet, and being prepared for anything.

Disaster Survival

Leaving the obvious for last, a DRaaS system gives you added assurance that your organization can and will survive a disaster. Without DRaaS, you might still be on the sidelines, gazing longingly at those bigger organizations that can afford DR — the old school kind, anyway.

Notes

Notes

Use DRaaS to prepare your business for the unexpected

In today's always-connected world, you can't afford downtime. Even when the unexpected occurs, your business or organization needs to be always on.

That's where disaster recovery as a service, or DRaaS, comes in. With DRaaS, the hard work of "what if?" planning can be passed to a service provider, meaning your team doesn't have to do the laborious work usually associated with the task. Even small businesses can be prepared for the unexpected with DRaaS. This book shows you how.

- *Plan for your specific needs — Assess the critical processes that you need to maintain, set recovery targets, and work with cloud providers*

- *Select a service provider — Learn the key factors involved in making the right decision for your business*

- *Ensure continuity — Build a disaster recovery (or DR) plan that ensures your organization survives a disaster*

Go to Dummies.com
for more!

FOR
DUMMIES
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.