InterVision

THE ULTIMATE GUIDE TO

DISASTER RECOVERY AS A SERVICE

Considering Disaster Recovery as a Service? This guide provides a business-focused perspective on IT recovery and how to leverage DRaaS to meet your availability requirements.

CONTENTS

CHAPTER

1

WHAT IS DISASTER RECOVERY AS A SERVICE?

Learn how the DRaaS process can help provide IT availability for your business

SAME TERM – DIFFERENT MEANINGS

With businesses relying more and more on technology for their critical operations, data protection and systems availability have become a concern that is extending beyond the IT team and into the executive boardroom. To address these concerns, many providers are developing new solutions, technologies and services promising to make IT Disaster Recovery easier and less expensive. With such a wide range of options labeled as Disaster Recovery as a Service (DRaaS), it can be difficult to understand what exactly the term means and how it can support your business needs.

"Without a robust DR infrastructure, a fire, a flood, an earthquake or even a really bad storm could not only take a business offline, it could take it out completely. Organizations increasingly rely on digital technologies for basic operations, and if the IT infrastructure is destroyed and data is permanently lost, it may be impossible to restore operations."

– Doug Hazelman: Vice President, Product Strategy, Veeam

DEFINING THE COMPONENTS OF DRAAS

At the most basic level, DRaaS is the process of replicating and hosting servers and their data via a third-party provider with the purpose of enabling a failover option in the case of an unexpected interruption or event.

To help create a more complete picture of the DRaaS process, let's take a closer look at its major components:

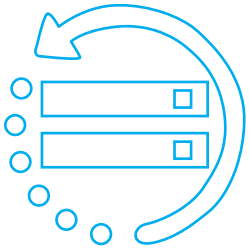

REPLICATION

Replication is the process of duplicating data from a production environment and transmitting it to the third-party host. Since most organizations have not fully virtualized their infrastructure, the replication technology should have the capability to frequently capture both virtualized and non-virtualized (physical) servers. This is often referred to as a "hybrid" solution. In instances where data is constantly added and changed, frequent data snapshots are essential in order to prevent data loss during failover.

FAILOVER

Failover is the transitioning of end-user access during an event to the third-party hosting environment. A DRaaS solution should provide a simple, straightforward process for declaring and failing over to the replicated environment. The speed in which the solution can "stand up" the replicated environment corresponds directly to the impact on business operations. Integrated tools and support should be in place so that failover can occur, even if your IT is unavailable.

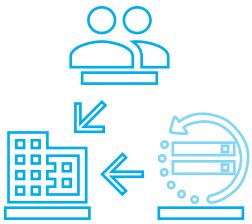

FAILBACK
Failback is the process of transitioning user access from the third-party failover site back to the original production environment. Once the impacted environment is fully up and running again, the DRaaS solution should have processes and tools in place that allow for a seamless switch back to production. Immediately following failback, replication processes should be re-engaged to ensure continuous protection.

IT RESILIENCY WITH DRAAS

When DRaaS is implemented correctly, it can provide more than just an "insurance policy" for your technology, because end users never experience the "disaster" in the first place. Active replication, along with a well-documented failover and failback process, delivers resiliency for your IT systems, providing the always-on user experience that customers and employees expect from modern business. When the cost savings of the cloud are leveraged, DRaaS is significantly less expensive and resource-intensive than fully replicating your existing production environment.

"IT resiliency / business continuity is a must in today's environment. The baseline assumption should be that IT components (i.e. servers, operating systems, networks, etc.) will fail at some point. And that the business functions that those systems support should continue to be supported and maintained."

– Allan Leinwand: Chief Technology Officer, ServiceNow

At InterVision, we understand that your critical systems, applications and business files are some of your most important assets. Seamlessly bridge any interruption with our tailored Disaster Recovery as a Service (DRaaS). So your business stays on, no matter what.

CHAPTER

2

THREE TYPES OF DRAAS

Understand your options of recovery & availability

DRAAS SERVICE EXPERIENCE TYPES

On a basic level, there are three types of DRaaS: self-service, assisted and managed. While there might be small variations in these models, DRaaS providers typically fall into one of these three options.

1. **Self-Service DRaaS**
   You get the tools to assemble your DR plan yourself
2. **Assisted DRaaS**
   You get the tools to assemble your DR plan yourself with DRaaS experts available for advice & assistance
3. **Managed DRaaS**
   DRaaS experts assemble your DR plan and manage all maintenance

SELF-SERVICE DRAAS

In this model, you will receive tools to perform DR backups and replication on your own. This means you're responsible for monitoring the status of your recovery environments and deciding how often you test that each solution is working correctly.

In a self-service model of DRaaS, your business is responsible for every aspect of the recovery planning, testing and management process. This means that when a disaster strikes, your IT team is completely responsible for executing the recovery.

Things to Consider:

- Typically, self-service models offer the lowest investment option, with the trade-off of time and resources to manage them
- Self-service DRaaS is best suited for organizations with internal DR expertise and IT bandwidth to manage the recovery environment
- A self-service solution can present challenges in scenarios where the IT team is not available or capable of getting on-site, like in the case of a natural disaster

ASSISTED DRAAS

Few DRaaS providers offer this option. In the assisted model, the DRaaS provider acts as an advisor as you implement, test and manage your solutions. You are responsible for all of the aspects of your company's DR plan, with the DRaaS provider standing by to assist if needed. If a disaster strikes and some IT team members aren't available, the DRaaS provider may step in to help with failover and failback.

Things to Consider:

- An assisted model may offer a lower investment option than managed DRaaS
- Assisted DRaaS is best suited for organizations that have IT resources, but are looking for a third-party expert to provide guidance and support during the implementation and management of a DRaaS solution
- Typically, assisted DRaaS will not offer a recovery SLA, meaning the client will ultimately be responsible for recovery during an event

MANAGED DRAAS

A managed model of DRaaS means the vendor is responsible for the planning, testing and management of your recovery strategy. A managed DRaaS provider should manage nearly all aspects of your recovery, so that your team can focus on other business priorites. In the event of a disaster, a managed DRaaS provider carries out your recovery process – whether this means working in conjunction with your IT team or performing the runbook operations alone.

Things to Consider:

- Managed DRaaS typically requires a higher investment than other DRaaS options but provides the highest level of support and assurance that the your recovery will meet your unique business requirements
- Managed DRaaS is best suited for organizations with high-demand IT teams that are involved in the strategic operations of their business
- Expect a recovery SLA matched to your unique requirements with a managed model to ensure your provider will be responsible for recovery in the case of an event

THREE TYPES OF DRAAS

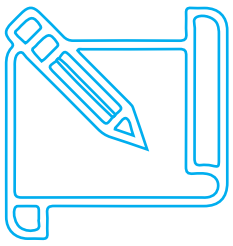

Self-Service DRaaS

Assisted DRaaS

Managed DRaaS

"Enlist the services of a DR expert, someone from outside the organization, to design and implement a DR solution or at least to audit what you have. The doer should not be the checker."

– Carl Mazzanti: Founder and CEO, eMazzanti Technologies

CHAPTER

3

DISASTER RECOVERY VS. BUSINESS CONTINUITY

What's the difference?

UNDERSTANDING THE ACRONYMS

It's common to see acronyms represent both disaster recovery (DR) and business continuity (BC). These terms, when used together as DR/BC, might suggest they are synonymous. Don't be fooled: their meanings are not the same.

People use the term DR/BC not because DR and BC share a common definition, but because they share a common goal of continued business operations. In fact, DR is just a smaller portion of BC.

*"Both approaches are key to business continuity. The first priority of a business continuity plan is to *avoid* disaster."*

—Matt Sprauge: Manager of Infrastructure Services, CDI Managed Services

THE PURPOSE OF DISASTER RECOVERY

The purpose of a DR plan is to recover your hardware, software and apps after a disaster. A BC plan, on the other hand, involves your finances, your personnel, your emergency plans and everything else that is a necessity to keep going and serving customers.

For this reason, it's best to have both DR and BC to achieve a fully functioning business. DR gives a specialized focus on the technology aspects of always-on business, while BC focuses on the broader aspects of staying available to engage customers and generate revenue. In simple terms, BC maintains overall responsibilities while a DR plan is being executed. Together, they keep business running smoothly and consistently.

"With so many businesses being IT driven, oftentimes business continuity cannot happen without IT having a solid DR plan. The biggest gap we run across is getting the business' workforce access again. Many IT departments are prepared to recover servers and networks, but have not considered how to communicate and re-enable the staff to resume work."

—Dustin Bolander: Chief Information Officer, Technology Pointe

WHAT SUCCESS DEPENDS UPON

Quite simply, the success of your business depends heavily on both DR and BC plans working in conjunction. Most of today's network outages occur not because of weather-related incidents, but because of human error. It's crucial that companies plan for the full gamut of events that could likely, and unlikely, impact their operations. For example, a structurally-fortified datacenter might guard against damage during a tornado, but it does nothing to prevent a cyber intrusion without virtual protection as well.

CHAPTER

4

UNDERSTANDING DRAAS, REPLICATION & BACKUPS

Know the difference to ensure complete protection of your data

OPTIONS THAT WORK TOGETHER

There are several technology options that exist with the goal of making data secure, available and recoverable. Knowing every solution type and the differences between them aids in good decision-making, flexibility for a tight budget and transformation for the future.

"If data is not backed up frequently, as in every 15 minutes or less, employees could lose a great deal of work, businesses could lose orders, all of which loses time, costs money and damages a company's reputation."

– Doug Hazelman: Vice President, Product Strategy, Veeam

DEFINITIONS AND DIFFERENCES

The most important differentiation between recovery options lies in the purposes each solution intends to achieve. On a high level, terms can be broken into two goals: retention or recovery. While some may blur the lines between these two goals, all will ultimately give emphasis to one over the other — and this emphasis is often determined by the data in question.

ARCHIVE

A collection of historical business records that are immutable and unchangeable, which must be retained for future reference should you ever need to produce evidence (perhaps when legal or regulatory questions arise). Archiving refers to the process of moving data that is no longer in active use to a separate storage device. Many organizations define archive data as any copies older than 120 days, whereas some may archive immediately when a matter is closed. Typically, archive data can be located online in to-disk format or offline in to-tape format.

ARCHIVE as a SERVICE (AaaS)

Also emphasizes maintaining a historical record, but unlike the "archive" definition above, a vendor is helping you with the act and storage of the archive data. It's typically assumed the storage is in a second datacenter managed by the vendor.

BACKUPS

A snapshot, point-in-time copy of your current on-premises data that you use to restore original file-level data if it's ever damaged or lost. This refers to the practice of copying computer files to another storage device, with additional copies being added to an external storage location as those files change. Backup data can be located online in to-disk format or offline in to-tape format.

BACKUP as a SERVICE (BaaS)

An approach where instead of performing backup on-premises, BaaS connects systems and data to a second datacenter, private, public or hybrid cloud managed by an outside provider. This outsourced method offloads the duties of managing tapes or hard disks and keeps data accessible or restorable from a remote location. Some people also refer to this form of data protection as Backup as a Service (BaaS), which is meant to place more emphasis on its purpose: to be able to recover using these backups.

IT DISASTER RECOVERY (DR)

A type of data protection that is usually implemented with the goal of returning systems to normal. Depending on the approach, recovery could be within hours, but yet typically takes much longer (days or weeks). It is typically accomplished using a combination of backups, SAN-to-SAN hardware-based replication, database replication or even software-based replication. With traditional IT-DR, the DR site requires the same amount of physical hardware and capacity as what is being used in the production or primary site, or this hardware must be procured before recovery can begin.

DISASTER RECOVERY as a SERVICE (DRaaS)

Emphasizes fast recovery from a technology disruption, restoring applications and networking often within minutes-to-hours so users can return to work again quickly. It is like BaaS, but instead of protecting individual files, it involves protecting the entire system or environment. On-premises datacenters or cloud-based datacenters are replicated to a second datacenter, private, public or hybrid cloud managed by an outside provider, removing the need for organizations to manage and maintain their own DR sites. It also eliminates the need for capacity planning because the service allows for scaling of hardware resources when a declaration occurs, reducing costs.

DISASTER RECOVERY on INFRASTRUCTURE as a SERVICE (IaaS)

References when an organization builds a recovery solution using replication or migration tools with an IaaS or cloud provider such as Azure, AWS or Google. In many cases, this solution is set up with the intent of reducing overall costs for DR, but the firm's IT team is responsible for the implementation, management, testing and recovery of the applications. It's important to note that failover and failback capabilities should be tested on a small dataset if you are protecting workloads from outside of that same cloud family, because, in some cases, the failback doesn't exist or requires a large effort.

PUTTING IT ALL TOGETHER

Oftentimes, an organization's budget has an impact on what solution is chosen, so the better you understand the importance of each dataset and application you need protected, the better you'll be able to stretch your budget and elect the best-fit solution. Also consider where your business wants to be down the road when choosing from these options. Selecting a BaaS provider that also offers DRaaS solutions, for example, may mean an easier transition later on.

THE HEAT SPECTRUM

To understand which recovery solution best fits your IT systems and applications, it's helpful to map them on a spectrum of hot and cold. The "hotter" a solution, the faster recovery capability it will have. For example, if a solution is labeled "hot," this means it can restore systems within minutes.

The "hottest" solution is not always the most cost-effective. That's why companies usually choose a solution that matches the needs of their applications to the costs and capabilities of recovery. With this in mind, a "cold" solution represents a slower recovery timeline. While the "coldest" solution may not be the fastest, it will often be the lowest investment.

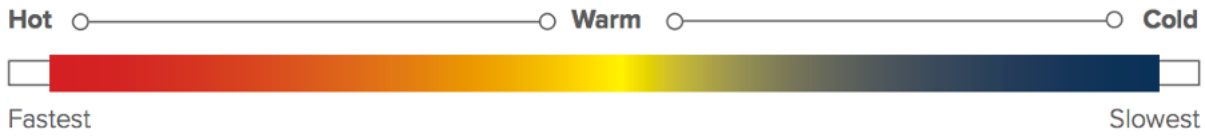

Archives are the "coldest" of all the recovery options, whereas **Backups** are next in line. If IT systems crash, backups can be retrieved to restore their applications again. Backups can be housed in the cloud or disconnected from running technology, which adds a layer of security to prevent the exposure of data. Data from backups that are housed in a physical form might take more time to identify during downtime. This is why they are considered a "cold" method of recovery.

Backups and archives may be done in accordance with a compliance mandate, as a precaution, against a cyber intrusion. In this scenario, they are usually stored off-site, most often in a vault.

For copies of data that are meant for shorter-term retention and typically for a faster recovery timeline, companies often use replication, which enters the IT Disaster Recovery (DR) realm of solutions. **Replication** often duplicates data as changes occur in your environment and provides access to recent iterations of your data during an event.

TWO CATEGORIES OF REPLICATION

Real-Time Replication:
Provides recovery in seconds or minutes. Real-time replication tracks and writes changes as they occur in an environment, so you can failback to an earlier version if needed.

Backup-Based Replication:
Backup-based replication records a full environment, then reports changes on a regular basis (typically once a day), based on the nature of the application.

WHAT "AS A SERVICE" REALLY MEANS

In each of the "as a Service" offerings, it's important to understand that three different service models exist within them:

- **Managed** – the provider is responsible for performing and managing the solution. In scenarios when the client's IT team is unavailable, the provider can even execute the recovery process on their own.
- **Assisted** – the provider gives the client expert advice on building a robust recovery strategy and setup, while also helping during a recovery event when needed.
- **Self-Service** – the client receives the necessary tools for the solution, but is on their own for everything else.

There are shared responsibilities whenever you are using the cloud. For this reason, know how much responsibility you want to take on, who owns what responsibility and which you share. Organizations should consider each of these "as a Service" service models with a critical eye to ensure proper alignment with needs.

CHAPTER

5

3 STRATEGIES FOR STRENGTHENING CYBERSECURITY WITH DRAAS

Understanding the role of DRaaS in security incident response

BUILDING A MORE RESILIENT BUSINESS

The modern-day client expects your business to be available no matter what disruption it faces. This fact, coupled with the evolving cybersecurity threat landscape, is pushing more and more businesses to adopt an IT resiliency strategy that integrates data security and recovery strategies together.

This means that security specialists and IT personnel must co-own resiliency tasks to acheive full mitgation — which requires a holistic approach.

"DRaaS and cybersecurity can go hand-in-hand to protect a company's data."

— Mike Smith, President of AeroComInc.com

1 INCORPORATE DR AS PART OF YOUR OVERALL SECURITY INCIDENT RESPONSE PLAN

It's important to formally consider security events as disasters and treat them with the same urgent attention. Why? Because a security event can have the same lasting impacts as any other form of disaster in terms of lost data, extended downtime and reputational damage.

According to a survey conducted by IDG Research, 46% of companies consider security incidents as "disasters" and 81% incorporate DR into their overall security plans. Taking this approach is the first step to building a holistic IT resiliency strategy.

Key Takeaway:
An experienced DRaaS provider will shore up protective services (firewalls, patching, etc.) for your company's recovery environment and incorporate cybersecurity scenarios into your DR testing and playbook documentation. Be sure to ask about the detailed response play for cybersecurity threats when vetting DRaaS providers.

"Firms need someone guiding them through best practices. Not basic things like antivirus, but bigger picture aspects like data retention policies, user awareness training and other information-related policies."
— Dustin Bolander, CIO of Technology Pointe

2 BALANCE PREVENTATIVE & RESTORATIVE MEASURES

Both preventative and restorative measures work in conjunction to protect data and secure information from unwarranted hands. This makes for an advantageous strategy to both IT professional and security specialist groups, in that you will always have a first line of defense, plus a Plan B.

As part of this complementary approach, some companies even do vulnerability/change monitoring in their recovery environments. Scanning your DR data for changes or corruption is a way to flag suspicious behavior, so you can take action in both DR and production environments if you identify a vulnerability.

Key Takeaway:
DRaas providers have your data, so why not have "a second pair of eyes" looking out for you? Vulnerability scanning in a DR environment does not impact resources in production and can potentially uncover threats that were missed initially.

3 PUT UP WALLS WITH THE 3, 2, 1 BACKUP STRATEGY

If you make it hard for cybercriminals to access your data, it will deter many from attacking you. One thing you can do to protect your applications and data sets from security incidents is walling them off from each other wherever possible. If cybercriminals gain access to one application, then they will have to jump over hurdles to gain access to further information.

A robust backup plan strengthens your data protection policy. To set up good walls for your data, a good place to begin is with the 3,2,1 strategy:

- 3 copies of data
- 2 copies stored locally
- 1 copy offsite

This strategy isn't an end-all-be-all because true resiliency demands vigilance and adaptability. That said, having multiple copies of your data dispersed in secure places gives you options during an event, upping the possibility of successful recovery from a crippling scenario.

HOW DRAAS IS A SOLUTION OF SECURITY/DR CONVERGENCE:

- Data replication and backup to a secure offsite location
- Testing and documentation to ensure protection practices are in sync
- Ongoing monitoring and encryption
- Team of experts for methodology, maintenance and recovery execution (ideal for overburdened IT teams)

BRINGING IT ALL TOGETHER

What does a successful holistic strategy look like? Let's review a sample use case of a ransomware attack.

Let's assume an attacker encrypts your IT systems with ransomware and demands payment. By using backups for long-term data storage and a replication solution for real-time changes in the cloud, DRaaS can offer checkpoint options for recovery.

Once an attack is recognized, replication should be paused immediately, so that the infected data doesn't spread into your DR environment. In this type of event you just invoke your DR plan to retrieve the most recent, clean copy of your data, then simply wipe your IT systems and reboot them with your DR copies. Your failover can be tested before changes are made in production to ensure there are no signs of infection. No need to pay an attacker for your data.

"The most prevalent security threat in the IT landscape today is ransomware. Frankly, the only and best way to protect against ransomware ties directly into disaster recovery."

– Clayton Hart, CEO at Diverse Technology Solutions

CHAPTER

6

DISASTER RECOVERY TESTING

An essential component to the ongoing IT Availability Lifecycle

ENSURING RECOVERY GOES AS PLANNED

Even with the best technology and process, unexpected problems can arise when executing a recovery plan – especially under the pressure of a disaster. One of the best ways to uncover and mitigate these issues in a controlled environment is to build a routine of comprehensive disaster recovery testing.

"DRaaS and cybersecurity can go hand-in-hand, to protect the most important aspect of a disaster recovery process implementation is testing and remediation. Chances are organizations will not get it perfect the first time. That's ok, what's more important is the ability to identify the issues and correct them, on a regular basis."

— Matt Sprauge: Manager of Infrastructure Services, CDI Managed Services

Testing also provides an excellent opportunity to train employees, perform maintenance and set expectations across the organization on recovery capabilities and limitations. For companies serving clients with “always-on” business expectations, testing is an essential component to the ongoing DRaaS Availability Lifecycle.

DRaaS Availability Lifecycle

EXAMPLE PHASES OF DR TESTING

Pre-testing

During the pre-testing process, a recovery team should review the entire plan to ensure testing runs as smoothly as possible. It’s important that each new recovery test challenges the technology and the team in ways that pervious tests have not. Additionally, efforts should be made to plan for scenarios where key IT team members may not be available to execute the plan.
Pre-testing activities may include the following:

- Reviewing recovery objectives and targets
- Reviewing process checklists and recovery runbook
- Verifying that production and the recovery environment are in-sync
- Understanding test objectives to set the right expectations
- Scripting out recovery for all VMs – to expedite recovery, especially in the event key team members are unavailable

Testing

If the pre-testing phase has been thorough, the testing process should be straightforward. During this time, systems and processes should be closely monitored for unexpected changes or problems.

Testing activities may include the following:

- Recovery of all physical and virtual machines
- Complete review of network connectivity
- Ensure production replication is not impacted by the testing

Debriefing

After the test is completed, the results should be reviewed and matched to the original test objectives. If issues are uncovered during the testing process, the recovery runbook should be updated to reflect the required changes.

Debriefing activities may include the following:

- Matching results to documented recovery objectives
- Discussion of how the process could have been streamlined for a faster recovery
- Communication of results to key stakeholders
- Documentation of uncovered issues in runbook
- At a high level, document objectives for the next test to prepare for and ensure a different testing scenario

EXAMPLE OF ISSUES AND CHALLENGES UNCOVERED DURING TESTS

- Network configurations: Networking settings may not account for the transition from production to the recovery environment
- Load balancing: Performance issues may arise if the recovery environment is not prepared to accommodate the load balancing requirements of production
- Firewalls: The recovery environment should mirror production to prevent access from unauthorized users, while also remaining available for authorized users
- Technology co-dependencies: Issues may arrise due to changing IPs, remote or proprietary technologies and lack of proper change management

CHOOSING A PARTNER

With the daily pressures put on IT teams to meet increasing business demands, committing to ongoing disaster recovery testing can be challenging. Unfortunately, statistics show that most organizations fail to pass their own tests. To stay ahead, many organizations are turning to a partner to help them manage their availability with Disaster Recovery as a Service (DRaaS).

At InterVision, we pride ourselves on delivering confidence in IT systems availability through our Recovery Assurance™ program. We manage the recover testing processes end-to-end so that your team can focus on other more-immediate priorities. Contact us today if you'd like to learn more about our specific approach to recovery testing and how we help over-burdened IT teams meet availability demands through Recovery Assurance.

CHAPTER

7

4 STEPS TO A PROACTIVE AVAILABILITY APPROACH

How to achieve disaster avoidance

DISASTER AVOIDANCE

You shouldn't need to scramble after a disaster strikes. For every minute of downtime, your business is losing revenue and damaging its reputation – which makes reactionary measures no longer sufficient for a company's livelihood. For this reason, it's best to approach disaster recovery (DR) from a perspective of "disaster avoidance."

To make recovering IT systems a cinch, a proactive approach needs to take precedence. As part of this goal, a new term has emerged: IT Availability.

Disaster Recovery Focus

- Invest in an "insurance policy"
- React to downtime and events
- Rely on backups to store data
- Treat DR and security separately
- Recover in hours to days
- Emphasis on technical infrastructure
- Develop minimal process and reporting
- Emphasis on avoiding "catastrophes"

IT Availability Focus

- Invest in ability to serve clients
- Be proactive to minimize risk
- Failover and failback to ensure service
- Secure recovery to protect data
- Recover in minutes to hours
- Emphasis on serving end users
- Process-driven and documented
- Emphasis on continuous improvement

IT Availability recognizes the co-dependencies that business has with its technology, whereas DR views technology in a boxed perspective. When companies fall on the side of reactive attention, often there are cultural frameworks that prevent proactive resiliency. The key is to change this way of thinking.

HOW TO ACHIEVE A PROACTIVE AVAILABILITY APPROACH

The best way to set the right culture for refining and improving preparedness is to open a dialogue with stakeholders. With input from these individuals and/or business units, follow these important steps:

1. Business Impact Analysis
Identify your critical business processes and functions, then define the recovery requirements and map technology dependencies.

2. Assess Risk of Impact
Identify risk scenarios, then determine the probability and impact of each scenario.

3. Develop a Strategy
Develop plans that address each scenario to mitigate risks.

4. Test the Plan and Update
Test your strategy frequently and update it continuously.

"More so than disaster recovery, business continuity is all about anticipating the faults, failures, and attacks that have the potential to impact the key services. It requires a careful and comprehensive Business Impact Analysis and the accompanying operating procedures to ensure that there is minimal disruption to the service."
— Allan Leinwand: Chief Technology Officer, ServiceNow

DRAAS AVAILABILITY LIFECYCLE

In addition to the four steps outlined above for proactive DR, it's essential to constantly improve your DR strategy so that it doesn't become outdated. One way to ensure this constant improvement is with the DRaaS Availability Lifecycle:

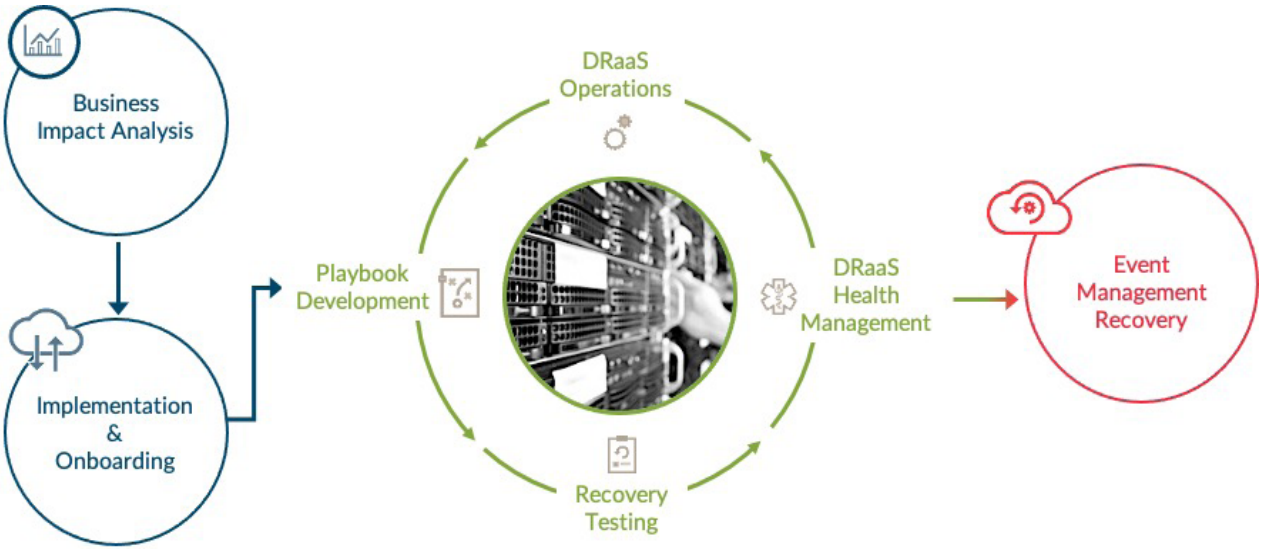

Business Impact Analysis The "Business Impact Analysis" stage takes note of how lost data and compromised applications impact overall business operations and revenue generation. Synthesizing this information, IT teams work with a DRaaS provider to tier applications into groups by order of importance.

Implementation & Onboarding During the "Implementation and Onboarding" stage, companies select the right DRaaS solution for each tier, accounting for required service level agreements (SLAs), technology capabilities and budget. A DRaaS provider gives advice based upon potential risk scenarios and assists in the architectural design and networking details. After the onboarding process has been completed, copies of data and applications are replicated into their chosen recovery environments for safekeeping.

Playbook Development In the "Playbook Development" stage, also known as the runbook development stage, IT teams consult with the DRaaS provider to document the order of operations for recovery, customizing steps for every scenario under the sun and noting each team member's responsibilities. The more comprehensive, the better. The DRaaS provider helps map out which operations need to occur before systems are returned to end users.

DRaaS Operations In the "DRaaS Operations" stage, the DRaaS provider focuses on the daily maintenance that ensures healthy DR environments remain healthy. This includes environment changes and resource monitoring.

DRaaS Health Management Similarly, the "DRaaS Health Management" stage emphasizes real-time monitoring, but with a focus on security, budget and compliance settings.

Recovery Testing Then comes the "Recovery Testing" stage. DR testing should be always on-demand and comprehensive. In some instances, tests should invoke your full IT team and DRaaS provider, and cut crucial members out to simulate a real-life scenario. Host pre- and post-event consultations to determine updates to the recovery runbook/playbook.

Recovery Event Management The "Playbook Development," "DRaaS Operations," "DRaaS Health Management" and "Recovery Testing" stages repeat themselves indefinitely, until the event of a real-life disaster, at which time the recovery process is executed in the "Recovery Event Management" stage.

CHAPTER

8

BUILDING THE BUSINESS CASE FOR DRAAS

DISASTER RECOVERY AS A SERVICE IS NOW MAINSTREAM

To provide continuous IT operations and protect valuable data, more businesses are turning to Disaster Recovery as a Service (DRaaS). Indeed, wider adoption of DRaaS can be largely attributed to the proven viability of cloud-based solutions; thus, the size of the DRaaS market will eclipse traditional DR solutions.

The following information provides tips on how to proactively build a business case for DRaaS, before your business is impacted by downtime.

MEETING KEY STAKEHOLDER DEMANDS

In the modern business, IT availability impacts the your entire organization and your ability to serve clients. Understanding stakeholder demands can help frame your conversation and approach to building a business case for DRaaS.

Technical Teams: Deliver IT systems availability and uptime

- Focus on operational efficiency
- Implement, operate and maintain technical solution and tools
- Drive coverage and compatibility of systems
- Execute tests, failover/failback

Executives and Board: Concerned with consistently meeting client demands

- Demand uptime and security of systems that drive service and revenue
- May lack deep understanding of risks and recovery capabilities
- Require confidence in protection of company's precious assets (data + revenue)

Business Units and Customers: Drive demands for application availability

- Create reliance on IT systems and services
- Core source of RTO and RPO requirements
- Direct knowledge of systems availability on projects, staff and revenue

Auditors, Insurers and Regulators: Need proof of uptime to meet specific standards

- Require recurring proof of disaster recoverability
- Demand significant documentation of controls and testing results
- Consume substantial staff time and focus when needs must be met

To receive equal participation from all parties, communicate how each role's demands will be met with your proposed DR strategy.

"IT must build the business case in the language that the "C" level management understands. IT understands the need and the importance of a solid disaster recovery plan. But, talking about data and backups will glaze the eyes of management. A better way to present a plan is explain how the plan will prevent the loss of business, profit and credibility."

—Karen Puchalsky: Founder, President and CEO,
Innovate E-Commerce

INVESTING IN AVAILABILITY, NOT AN INSURANCE POLICY

According to a survey conducted by InterVision, 60% of IT professionals consider protecting their business against technology-related disruptions "extremely important" but only 19% consider their organization's funding for IT disaster recovery (IT-DR) to be "very good." It's common for organizational leadership to see investments in IT-DR as "insurance," however an effective DRaaS solution should be focused on delivering availability for your customers and internal stakeholders.

> *"It is important to elevate the disaster recovery conversation to the business level. Begin by identifying the business process we are protecting and the underlying technology that enables that process. For the technology — the data, applications, and systems — evaluate the options for recovering it in a quick and cost-effective manner. Select the recovery option, implement it, and do regular drills to ensure the process and implementation is functional and sustainable."*
>
> — J Wolfgang Goerlich: Director of Security Strategy, CBI

Considering DRaaS as an investment in the ability to continuously serve clients, secure data and proactively improve operations can help open the dialogue beyond the traditional disaster recovery focus.

THE FINANCIAL IMPACT OF DOWNTIME

Even with increased customer demands for uptime, many organizations struggle to make the business case for implementing an effective DR solution before downtime impacts them directly. With market research indicating that just one minute of downtime costs businesses an ever-increasing amount, a proactive approach to disaster can mean substantial cost savings.

To accurately understand the impact downtime can have on your business, it's important to consider direct "hard" costs as well as "soft" indirect costs.

"A business case for DR and resiliency is simple: the impact of even a small loss of service is devastating to your brand(s) and the company as a whole. It cannot be measured simply in lost transactions. Loss must be measured with a social network multiplier."

– Adam John: President, Sterling Solutions

Direct Costs

- Lost revenue
- Breach of customer contracts or service level agreements
- Lost inventory or supplies
- Non-compliance fines

Indirect Costs

- Reputational damage
- Brand impact
- Negative publicity
- Loss of competitive advantage

YOUR ORGANIZATION'S OPTIONS TO EVOLVE

Based on the types of products or services your organization provides, your clients and stakeholders will have different IT availability expectations. However, as reliance on technology continues to grow, you can logically expect those expectations and requirements to grow in parallel.

IT AVAILABILITY MODEL

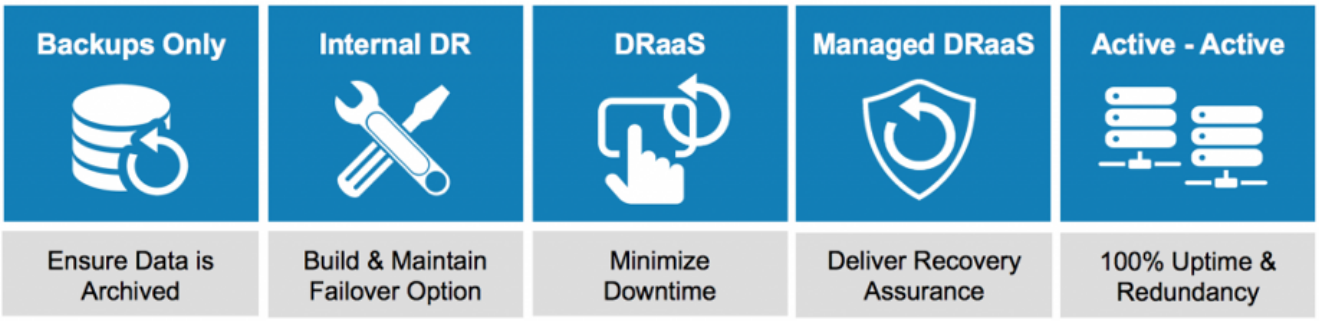

Understanding where your organization falls within this "IT Availability Model" can help outline options as your IT requirements grow. It's important to understand not every business needs to maintain a fully active-active IT environment. Consider matching your business goals to your budget when selecting the right technology solution.

CHAPTER

9

HOW TO COMPARE FAILOVER TARGETS FOR DRAAS

Choosing the right fit for your recovery

DRAAS TARGET OPTIONS

As the modern business landscape moves further to cloud, it's a key time to consider whether cloud is also a good fit for your IT disaster recovery (IT-DR) plan as well. Disaster Recovery as a Service (DRaaS) is essentially cloud-based DR, able to run in any type of cloud architecture or datacenter infrastructure.

1. On Premises – your own datacenter on your own gear
2. Hosted Infrastructure – multi-tenant or private environment
3. Cloud – AWS, Azure, Google Cloud, etc.

WEIGHING YOUR BUSINESS FACTORS

When considering where to target your DR plan, InterVision has seen five primary drivers that dictate where companies ultimately place their DR:

1

If you have a cloud-first strategy but not much experience with cloud, most companies start with DR in the cloud first

Broader Cloud Strategy

InterVision has often seen organizations whose long-term strategy involves a move to the cloud that has been slowed or stalled in the process gain value from a cloud-based DR environment where IT teams can get their feet wet before moving critical workloads into a cloud production environment. Giving IT teams with little cloud experience exposure to a cloud DR environment with copies of your data is a lower risk than data in use in a primary production environment. If this is your scenario, you probably already know what cloud your business wants, or what your business wants to do in the cloud. This cloud environment will be your failover target; you then figure out a low-cost option in that target for your team to learn the ropes.

2

If you have essential physical applications that must be virtualized, most companies explore a hosted solution with hybrid options

Mixture of Physical & Virtualized IT Systems

It's common to have legacy physical systems that are critical to business operations. However, these older applications tend to hinder the rest of the business from moving fast into the future. Sometimes these legacy applications can't be controlled or retired. For example, airlines must first check the no-fly list before taking off, a federal application that is separate from the airline's internal systems. Don't let a scenario like this hold you back from success. With a hosted infrastructure environment, you can provision a hybrid setup for failover operations in the cloud, even if you have a few or more applications still running on-premise. A hosted DR target tends to have the most flexibility in levels of features and capabilities to protect assets. Plus, a hybrid environment for DR usually can meet pricing on par with the cost of cloud.

3

If handling DR internally isn't core to your business strategy, a trusted third party can own DR management on your behalf

Preferred Service Level & Model

Customers have long dictated the sway of the market. Now, there's a demand for personalized, fast service in an always-on model. Think about e-commerce, streaming and ridesharing apps. All have upended their relative industries and emphasize a continuous level of service. For this reason, companies are rethinking their approach to technology in order to better meet their customers' desires. If handling DR in-house isn't a core goal of your business, we've found that offloading this responsibility in full to a trusted third party gives an IT team the ability to drive the business forward. A third-party vendor should be able to cater to your needs in any failover target and handle all aspects of design, implementation, testing and maintenance. The question then becomes what level of recovery speed you need and the cost of achieving it.

4

If cost is a concern, your environment will dictate the best options. Consult with a third party for an assessment

Cost Components

Cost is always a factor in any business decision. If your business is unable to dedicate large sums to DR for the short-term and this is the most important decision criteria, it's key to think about where you eventually want to go before simply selecting the cheapest option out there. Maybe you won't be able to provision the best solution for DR right now, but you could find a solution for a reasonable price that's adjacent to where you want to go and that would enable a smooth transition down the road. We've found that cloud tends to offer the most options for a tight budget, with the ability to move workloads easily as needed, but the best fit will ultimately depend upon your environment.

5

If fast failback after a disaster is a priority, evaluate your target options with this as the key consideration

Failover & Failback Capabilities

There's no point in sending your critical data to a DR environment during an emergency if you'll never be able to get it back. A disaster event will never be fully resolved until workloads are back in their normal location again. If your business depends upon having normalcy again quickly, don't just consider the path your data will take during a disaster; consider how and when you'll send that data back to normal production. Some cloud environments, such as AWS, don't allow failback easily, so a managed service provider will be key. If having your data back in a timely fashion is important following a disaster, a hosted infrastructure might be a good option to consider in addition to a specific cloud target.

"IT resiliency / business continuity is a must in today's environment. The baseline assumption should be that IT components (i.e. servers, operating systems, networks, etc.) will fail at some point. And that the business functions that those systems support should continue to be supported and maintained."

– Allan Leinwand: Chief Technology Officer, ServiceNow

DRAAS STARTS WITH STRATEGY

What you care about most will ultimately point you in the direction you need to go for DR. Start with your goals first. Consider whether you want to reduce total IT costs, improve confidence in recoverability, gain flexibility in both full and partial failover capabilities, get out of the hardware business, or gain coverage across multiple datacenters. Knowing these goals will help you with challenges ahead and navigate the waters to success.

A strategic service provider (SSP) can assist your business in selecting the right failover target for your IT-DR plan. Sometimes organizations use DRaaS targeted to the cloud as a first step toward a larger cloud migration, demonstrating an early win for stakeholders for greater buy-in. Read our blog post, "Considering DRaaS as a First Step for AWS Cloud Migration," for more information.

CHAPTER

10

9 TIPS TO A SUCCESSFUL DRAAS IMPLEMENTATION

How to create a disaster recovery plan with confidence

MEET BUSINESS REQUIREMENTS AND PROTECT DATA

Disaster Recovery as a Service (DRaaS) can be an effective solution to combat downtime and achieve IT availability. If approached in the right way, it can provide comprehensive proof of recovery to stakeholders and, most importantly, increase your availability during an event.

"In the event of a malicious attack, a company should have systems in place to keep operational or at least backups where the company is not affected or very slightly affected. In the event of a total disruption of the business, it is too late to mitigate and you will likely see dramatic costs to the business, especially small or mid-sized businesses. Being proactive rather than reactive is the key."

— Braden Perry: Partner, Kennyhertz Perry, LLC

TIPS TO IMPLEMENT A SUCCESS DRAAS STRATEGY

1

Disaster Recovery is more than insurance

IT Disaster Recovery (DR) is an investment to protect your business. The reality is that interruptions to your business will happen and it's best to have a plan to minimize the impact. In an age where being offline is unacceptable, a proactive approach is needed to keep downtime from occurring in the first place. That is why it's important to view recovery from the perspective of IT availability.

2

Gain company-wide investment

Having all stakeholders (executives, directors, managers, board members, etc.) equally invested in your recovery strategy will ensure you receive the proper resources. Aim to have a company-wide agreement that a proactive approach to your company's IT systems is necessary to ensure continuous operations. With this level of commitment, you will always have the resources you need to recover, plus prioritized attention on availability.

DR has evolved from a view of insurance to a corporate priority. However, a recent survey shows those removed from the implementation of DR are overconfident in their IT team's ability to execute. To answer the disconnect, this white paper examines the motivations behind adopting Disaster Recovery as a Service (DRaaS) to protect data assets and empower business objectives.

3

Cover all applications at the right level

Sometimes companies over-invest in technology, replicating everything and making their recovery more expensive than needed. Others do the opposite and are more vulnerable to downtime by under-investing.

DR is never a one-size-fits-all approach. The goal is to assess each application and match it to the right recovery technology. Taking a tiered approach ensures you receive the right recovery level for your budget.

Organizing applications into tiers of recovery by their level of urgency enables the right focus during an event. With this goal in mind, Recovery Waves go a step further. They define the order in which those applications, within each specific tier should return to service. Learn more about our unique "Recovery Waves" process.

4

Combine high-availability with DR

High availability technology is great for instantaneous RTOs on local issues like hardware failures or data corruption, but these technologies often need to be physically near each other to work. To plan ahead for non-local issues, pair your high availability technology with your DR plan to ensure continuous service. To cover the full range of disruptions, have the two working together.

5

Plan for multiple causes of downtime

Many organizations tend to think DR planning is protecting your business against environmental disasters when in reality, environmental causes are among the least likely causes of downtime.

TIPS TO IMPLEMENT A SUCCESSFUL DRAAS STRATEGY

5 (cont.)

Plan for multiple causes of downtime

Avoid leaving your business vulnerable by planning for a wide range of potential events. After all, the purpose of DR is to plan for the recovery of your data and applications, no matter what the scenario. For example, you don't want to spend all of your resources building a tornado-proof datacenter only to have your data corrupted by a security breach.

Common interruptions:

Mechanical: Power Outage, Hardware Failure, Networkf Failure

Environmental: Flood, Earthquake, Fire

Human Causes: Unauthorized Access, Human Error, Cyber-attacks

6

Test regularly & differently

For confidence in recovery, you should know how your DR environment will perform during an event. It's important to test often and for different scenarios. Be sure to plan for instances where your IT team will not be available.

Testing once a year doesn't take into account the changes that might occur to your production environment, which can affect your ability to recover at any given point in time. The more often you test, the more you'll understand how your DR plan will perform. Listen to Developing a DR Runbook podcast for more tips on effective DR testing.

7

Leverage all of your resources

If you have a datacenter in a separate regional location from your business, why not use it to store your replicated data? Likewise, if you are running tape backups, it's best to use them to complement your DRaaS solution for long-term data retention. This ensures complete recovery of your data and efficient use of your budget. Similarly, pairing your cloud strategy and resiliency program together can reduce overall costs while ensuring continuity of operations.

TIPS TO IMPLEMENT A SUCCESSFUL DRAAS STRATEGY

8

Ensure proof for constituents

If your company is receiving pressure to prove your DR plan to constituents—whether it's auditors, regulators, clients, board members or investors—it can be tough to give them the evidence they need. Oftentimes, these stakeholders are looking for unbiased, third-party proof.

A DRaaS provider has the ability to offer tangible, third-party validation of a DR strategy. When you talk with a vendor, be sure to ask what this evidence will look like and how they will deliver it. A good DRaaS provider will provide several types of evidence (such as signed certificates of testing, a playbook and a dashboard for real-time analytics).

9

Pick the right DRaaS provider

When selecting a DRaaS provider, it's important to consider not only their recovery technologies, but also their operational maturity and expertise. Look for a provider with both the technology you need and a culture that's matched to your business objectives.

InterVision's flexible and proactive approach to DRaaS ensures organizations receive complete confidence in their systems availability and data protection so IT teams can focus on other, more pressing business initiatives. We can deliver our DRaaS and BaaS solutions using any cloud location, such as AWS and Azure.