
The Essential Guide to Cloud-based Backup and Disaster Recovery

5 Steps to get started quickly

veeam

CSP PARTNER
Platinum

2017
veeam
CLOUD SERVICE PROVIDER
OF THE YEAR

iland[®]

When disaster strikes, whether large or small, your job is to keep things up and running. You need to build a plan of action. So, where do you start? You may need a backup solution, you may need a disaster recovery solution or you may need a bit of both.

Developing your company's business continuity plan begins with understanding the technologies available and matching them to the use cases you have. Then, after selecting a vendor, you can implement the solution, document your emergency plan and test the entire thing – so you can confidently pull the trigger when the time comes.

In this guide, we'll go through those steps, one by one, in an approachable, action-oriented way. iland, named by both Gartner and Forrester as a leader in disaster-recovery-as-a-service, has been helping companies like yours solve the business continuity challenge for over a decade. Our team of disaster recovery experts work with customers every day to implement cloud backup and disaster recovery solutions. We can go far beyond simple backup to ensure all your key workloads across virtual and physical environments are protected – and the disaster recovery process is tuned to your business priorities and compliance needs.

So, if you're looking for assistance to put your business continuity plan in place – reach out to us at iland.com/contact, anytime.

Now, let's get started.

Step 1

Understanding the Technology Options

BACKUPS

Backups are copies of your virtual and physical systems, including all the data in those systems.

Backups

Backups are copies of your virtual and physical systems, including all the data in those systems.

- > Backups are typically done daily, and stored in increments beyond that. Companies often adopt a retention schedule such as keeping daily backups for a month, monthly backups for a year, and quarterly backups beyond that.
- > Backups are often mandated by operational or regulatory requirements.
- > Backups can be stored on-premises, traditionally on tapes or archival disk – or, nowadays, can be shipped to the cloud. The cloud-based options provide:
 - > Geographic diversity from your main site
 - > Low cost options
 - > Straightforward management

While backups provide an easy way to ensure data is preserved, they have weaknesses. Backups are a poor solution to time-sensitive system recovery because each VM has to be manually restored before it can be accessed. This can be very time consuming. Furthermore, because they are taken daily, backups do not provide a very current view of the system (see sidebar). Restoring from a backup could mean the loss of 23 hours of data.

DRAAS

Disaster recovery is a system of replication designed to minimize downtime. It creates a copy of the VM at a secondary location, and can fail-over in seconds or minutes.

Disaster Recovery

Disaster recovery is a system of replication designed to minimize downtime. It creates a copy of the VM at a secondary location, and can fail-over in seconds or minutes.

- > Most DR systems leverage a continuous replication technology, which can record second-by-second changes on your primary system or a snapshot-based replication which takes copies in 15-minute increments. These functions drive the Recovery Point of the system – from seconds or minutes (continuous) to minutes or hours (snapshot), depending on the configuration.

- > While DR can be set up between multiple sites within a company, a cloud-based option ensures:
 - > Geographic diversity, in the case of a natural disaster or power outage
 - > Lower costs, eliminating the cost of a secondary data center and infrastructure
 - > The support of an engaged 3rd party to help in an emergency

Replication is the preferred method for minimizing downtime because the pace of replication minimizes data loss by keeping the secondary copy up to date, creating a very low Recovery Point. Secondly, the mechanism of replication, at the host-level, syncs multiple VMs at the same time without impacting performance.

However, most replication systems do not create or store long term backups as part of the replication function – and those backups might be required by your operations or compliance teams.

“For systems where uptime is critical to business operations and data is rapidly changing, a disaster recovery solution is the right choice for minimizing lost revenue and operational damage.”

Are backups the right solution for you?

If you have regulatory or operational requirements to keep backups, they are, of course, necessary. Furthermore, they may be sufficient as a business continuity tool if the application is non-critical and the data is not rapidly changing. For example, a billing tool used monthly to generate invoices could be down for 2 weeks with limited impact.

Is a disaster recovery solution the right solution for you?

For systems where uptime is critical to business operations and data is rapidly changing, a disaster recovery solution is the right choice for minimizing lost revenue and operational damage. Most systems even enable you to tier the priority of workloads, ensuring the most critical are brought up first, minimizing the business impact of a disaster.

Do I need both?

This is the case for many organizations. Many protect critical workloads with a disaster recovery solution and use backups to meet compliance requirements. Often, the on-demand pricing of cloud services means it can be cost effective to just protect it all, knowing much of the cost is incurred only in a disaster, minimizing lost revenue and operational damage. Most systems even enable you to tier the priority of workloads, ensuring the most critical are brought up first, minimizing the business impact of a disaster.

“A skilled cloud provider will have familiarity with both DR and Backup – and will not encourage you to make do with an incomplete solution .”

When seeking a joint solution, consider:

- > You may want to find a single vendor to handle both – having one cloud partner for business continuity is often easier for the organization
- > There may be integrated technologies that can handle both backup-as-a-service and DRaaS from a single on-premises installation.
- > You'll definitely want proof – an audit trail – for both DR and Backups, if you anticipate external oversight of your business continuity efforts

A skilled cloud provider will have familiarity with both DR and Backup – and will not encourage you to make do with an incomplete solution. Work with someone with tenure, positive analyst ratings, and a commitment to customer service.

Backup and DRaaS – an RTO and RPO example

Your systems are backed up every night at 7pm. You also are running a DR system with continuous replication. Disaster Strikes! At 6am, a power outage hits the building, taking down your data center. How do you recover?

DRaaS Scenario

- > At 6:01am, you press “Fail Over” on the DR solution.
- > Within 4 minutes, at 6:05am, the systems are back up in the cloud. The Recovery Time (RTO) is 4 minutes.
- > The data is current as of 6:00 – when the power went out. So, your Recovery Point (RPO) is 5 minutes.

Backup Scenario

- > At 6:01am, you go find your backups.
- > Then, you spend the next few hours finding, fixing and re-starting your infrastructure (assuming your power is back on).
- > At 5:00pm, a few of your major systems are back online. Your RTO on those systems is 11 hours.
- > The systems are restored to the last backup time – which is 7pm yesterday. So, your RPO on those few systems is 22 hours.
- > The remaining systems take longer, and thus have much longer RTO & RPOs.

Step 2

Categorizing your IT systems

The next step in developing your business continuity solution will be to categorize your own IT environment to identify which workloads need protecting, which are business critical, and which may be less important to ongoing operations.

The next step in developing your business continuity solution will be to categorize your own IT environment to identify which workloads need protecting, which are business critical, and which may be less important to ongoing operations.

STEP 2: Categorizing your IT systems

Disaster events can be large or small – and while most businesses will not face a wholesale disaster that touches their entire infrastructure, it is impossible to predict where a disaster will strike. So, it's best to take a holistic approach by starting with a pretty good inventory of your IT systems.

“...it is impossible to predict where a disaster will strike. So, it's best to take a holistic approach by starting with a good inventory of your IT systems .”

Inventories and Notes

The first step to identifying the needs of your IT environment is getting a solid list of potentially impacted systems. Talk to your IT department, stakeholders from other business units, and even your security team for input. You're looking for:

- > On-premises systems that support business applications – these tend to be long-running and can even be physical (not virtualized) systems
- > Anything running in your virtual environment on-premises
- > Systems housing departmental applications, which may be outside your data center but are still on-premises

- > Systems used for development and testing, which are typically transient
- > Applications housed in the cloud
- > SaaS applications which are consumed entirely via the internet

Typical organizations can have up to a few hundred applications, though some rare over-achievers top 1000. The idea isn't to necessarily categorize every one, but rather cast a wide enough net to ensure you're capturing the ones that are business critical. At iland, we help our customers identify essential workloads for backup and DR protection with a tool called iland Catalyst which provides full visibility into your existing environment so that you can select various workloads, understand performance impacts and resource needs in the cloud and simplify your migration to cloud backup and DR.

For business critical workloads, both on-premises and cloud-based, it's important to consider a few factors:

1. How frequently is this system accessed?
2. How often does the data in the system change?
3. How hard is it to recreate the data – or does it have value at all?
4. Is it interdependent with any other systems?

Access Frequency drives RTO

The Recovery Time Objective is the time between when the system goes down and when it comes back up again. Sometimes, that's just the time it takes to reboot a system, which can be measured in minutes. Other times, it takes a great deal more to revive a dead system – depending both on its manner of dying and on the business continuity plan you have in place.

If a system is used infrequently – for example, if it is the monthly invoice system – then for 3 weeks of the month, it's untouched. It could go down for 48 hours, and no one would be the wiser – or at least materially impacted.

On the flip side, if the system is your public-facing web site, from which you take an

order every 5 minutes, 30 minutes of downtime can be extremely noticeable and expensive to the business.

For infrequently-used systems, you can be more lenient on RTO targets. Maybe recovering from a backup – which often takes hours or days – is sufficient. For those that are highly-accessed, you'll need something that measures recovery in minutes, and thus, doesn't begin with "procure new hardware."

“While all data is typically treated as gold by an enterprise, some data is more important than others.”

Data Update Frequency Drives RPO

The Recovery Point Objective is a measure of how much data you lose while the system is offline. Some continuity tools take a snapshot of the system daily, while others do continuous replication. Daily or hourly snapshots might be sufficient – or they might be throwing out gobs of data that will later have to be recovered (or forever lost).

If the system is updated infrequently – say, once a week – then you can be offline for 24 hours and the last known copy of the system is still accurate. A snapshot you took last night of a system that won't change until Thursday is perfectly good.

However, if the system is updated frequently, you could lose a great deal of data. Weather prediction models tend to be updated twice daily, so, if your system is out for 3 days – and you revert to the last known daily snapshot - you've lost 6 updates. If it's updated more frequently (like sales data off your web site), you're losing a great deal more still. But, this brings us to the next question...

How replaceable is the data – really?

While all data is typically treated as gold by an enterprise, some data is more important than others. Sales data tends to be critical, along with order information, fulfillment, accounting data and so on.

But, in the example above, what good was the weather model from 24 hours ago? Does it have any useful value now that a new, more current, model is available? Also, those days have passed – and it rained (or didn't) – regardless of what the model said. So, really, that data probably isn't critical at all. Losing 3 days of it simply means waiting until the next available model is released.

Some data is needed long past its created for legal or operational reasons. Other data is naturally ephemeral, and requires no further storage. Identifying which systems hold which data is critical to accurately tuning your investment in continuity tools.

Interdependence – the final straw

As a final step, identify which systems are actually interdependent. Why? Because sometimes an innocuous and uninteresting system is actually the key to the operations of a far more critical system. So, always double-check dependencies to ensure you're properly classifying all the machines.

Bringing it all together

Now, you should have a pretty good list of systems and their characteristics. You can stack rank them in terms of business criticality to remain online (RTO) and data change (RPO), and begin to map them to the technologies available in section 1. You will likely find that:

- > The most critical systems to your business should be protected with a proper continuous or frequently-replicating Disaster Recovery solution
- > A second-tier of less critical but still core systems should also be protected with a DR solution, though they may be able to wait their turn to be initialized in a disaster

- > A third tier of systems might be fine with simple backups, with the understanding that days or weeks may pass before they become available again
- > A final tier – like testing systems – may be entirely transient and thus they could be lost with little downside

“The most critical systems to your business should be protected with a proper continuous or frequently-replicating Disaster Recovery solution.”

Before you set off to implement this plan, bear in mind a few things:

1. You'll want to communicate this plan to business stakeholders to get their opinion on the criticality of systems. After all, it's theirs to prioritize too.
2. Check out the prices from DR and Backup vendors. Particularly with cloud-based DR-as-a-Service, the prices could be low enough to make it easier and faster to just protect the bulk of your systems with the same solution.
3. Cloud-based workloads often come with the option for something like cloud-to-cloud recovery, which is effectively DRaaS for cloud systems. You can fail over to another location in the same cloud – or another cloud entirely.

Step 3

Implementation

When considering implementation of a business continuity solution, you need to consider how – and how much – you'd like to invest. Step 2 brought you a good assessment of what it will take to keep the business running. Now, it's time to pair it with your appetite for on-premises investment, capital and operating expenses, and ongoing management.

Backups

Many types of software can help you take regular backups. In virtual environments, there are built-in mechanisms, as well as options for backup management from a number of systems management and virtualization management vendors. Typically, backups are taken in regular increments and stored on a schedule – daily for a week, weekly for a month, monthly for a year. But, where are they stored? Guidelines suggest:

- > Have at least three copies of your data
- > Store the copies on two different media
- > Keep one backup copy offsite - and preferably well outside the impact zone of any local disaster in a safe, secure location

“So taking the backups is easy. But, storing them becomes a pain. ”

So, taking the backups is easy. But, storing them becomes a pain. Off-site storage can be ok if you happen to have a secondary data center, far enough from the first. For many organizations with a cluster of local sites, that far-flung secondary data center is out of the question. That's why many consider cloud backup options.

Cloud backups provide:

- > An easy-to-implement shuttling of your backups to a third-party location
- > Operational expenses in place of up-front infrastructure, data center, power, cooling and staffing costs

- > Geographic diversity within your region of governance – you often don't want to leave the country, or additional legal complexities could emerge
- > Very easy expansion options as your backups and systems grow in size
- > Ideally, a simple pricing structure

You'll want to consider:

- > Which provider will be best positioned to help configure and deploy your cloud-based backups – and support you in the event of an emergency
- > How much bandwidth you'll need to send your data over the wire – with whatever frequency you deem necessary
- > Whether the same cloud provider can support both backups and DRaaS - making vendor management more straightforward

Disaster Recovery

There are a number of technologies that can support continuous or frequent replication in a disaster recovery solution. To get recovery times and points within minutes, you'll need one of these solutions. The good news is that they are remarkably capable technologies that require minimal installation and configuration. The bad news is that like any business continuity solution, they need a target environment. Just as with backup, the secondary location is paramount. For those companies with two (or more) geographically-diverse data centers, there may be a plan for

“Just as with backup, the secondary location is paramount...often the most cost-effective solution is cloud-based DR, known as Disaster-Recovery-as-a-Service. ”

shunting replicas of data between those sites. However, you'll still need to maintain infrastructure capable of recovering those systems, idly sitting at the opposing location. Anyone without a secondary site, however, is out of luck.

The alternative option – and often the most cost-effective solution – is cloud-based DR, known as Disaster-Recovery-as-a-Service. An industry-leading DRaaS solution provides:

- > Continuous or rapid-enough replication to ensure you don't lose a lot of data
- > An enterprise-class target cloud environment
- > Strong post-failover management tools, to ensure you can seamlessly operate your workloads for the duration of their stay in the cloud
- > Integrated security features that mimic those you'd have on-premises
- > A straightforward manner of failing back, once the emergency is over
- > A clear, cost-effective pricing model that favors operational expenses, and on-demand pricing for resources in an emergency

In choosing a DRaaS solution, you'll want to consider:

- > What your network configuration will look like – and can the DRaaS provider accommodate it? Some configurations require more nuanced implementations
- > Whether you need support for physical systems or co-location
- > How much support you'll need from the provider – from architecture to configuration to testing and support during a disaster
- > Whether you and the DRaaS provider have adequate bandwidth available to replicate the data?

“Many organizations find themselves buying a mix of DRaaS and cloud-based backup solutions in service of their business continuity plan. The best approach is to find a cloud vendor you can trust to help craft the plan, architect the solution, and support you throughout the deployment and, as needed, a disaster.”

Many organizations find themselves buying a mix of DRaaS and cloud-based backup solutions in service of their business continuity plan. The best approach is to find a cloud vendor you can trust to help craft the plan, architect the solution, and support you throughout the deployment and, as needed, a disaster. But before you're done, you'll need to build the actual DR plan – and test it, which we we will discuss in Steps 4 and 5.

Step 4

Build the Plan

Once you have your DRaaS solution implemented, it is important to spend a bit of time thinking through the broader plan. Obviously, you've already worked out how your technology will fare, once you press the red button. But, the human element of declaring a disaster bears some consideration too.

What is a disaster?

You should consider outlining, either globally or by tier of system, what sorts of situations constitute a disaster for your company. Impending weather events of more than 50% likelihood of catastrophe? Hardware failures with more than 24 hours of guaranteed downtime? Thinking through these potential occurrences in advance will remove some of the burden of last-minute decision making.

Who can press the button?

Ideally, someone involved in crafting your DR plan would man the controls at the time of the disaster – but that can't always be guaranteed. Weather, logistics, staffing changes and other factors all can mean someone else might have to enact the plan. Make a list of folks who are authorized to make the call – and ensure they know what they've signed up to do.

What's the procedure?

Beyond failing over systems, disaster recovery plans will usually include communications to the company and executive team, interactions with your DRaaS provider, testing procedures to ensure a successful failover, and more. Create a simple, short document to walk whomever is at the helm through everything they need to know – and keep the acronyms and jargon out of it. No one has time to learn new lingo in the middle of a crisis.

Who can they call for help?

Ensure that with that procedure document is a list of names, titles and contact information for those who can help out – both internal to the company and at your DRaaS provider.

What if you've gone dark?

Assume the worst in building this plan – no internet, no power, no land lines. Ensure you have hard-copies of the plan in key places or on people's laptops and phones.

And, of course, your staff, procedures and technology are all subject to change, so revisit this plan periodically – perhaps during your annual/bi-annual DR test – to ensure it's all still valid and up to date.

Step 5

Test Your Disaster Plan

Things change in data centers and in life. New applications emerge. Old applications pass into history. New people join and new systems come online. And, while it would be nice to guarantee that with every shift, someone carefully considers the impact on the DR plan – but that's not usually the case.

The problem is, when disaster does strike, it is critical that you trust the system you have in place. If not, you might be reticent to press the button – rendering the whole exercise useless. In fact, study after study has shown that the biggest failure point in DR planning is the ultimate fear of enacting the plan – due to concerns about whether it will work. If the litany of systems changes looms large as the disaster bears down, it's easy to see why a company might abort.

The good news is that this can be easily avoided by testing the plan regularly. Some organizations do it on a monthly basis, while others on a quarterly basis. The rate of change of the systems in place should drive the frequency of your tests.

Most DR systems enable non-intrusive testing, which is invisible to the end users. They bring up the systems at the secondary site without taking down the primary site. But, some organizations prefer to do a full test – simulating a real outage by literally pulling the plug. You can be certain those companies are fully confident when the time comes. The two types of testing are called Failover Tests and Live Failover:

Failover Test

A failover test is a safe, non-disruptive option. This will bring up the VMs that have been replicated to your DR environment, and allow you to do whatever testing you need to do, without impacting anything in your production environment whatsoever. Once you end the test, the VMs will be removed, and the data will begin replicating again (some technologies continually synchronize in the background even while you're testing). Ideally, your DRaaS provider should let you test any time, for minimal cost.

Live Failover

This is the failover designed to simulate a full disaster. A live failover will assume that your production environment is down, and that the VMs being

failed over are now supposed to be the primary systems. Once you have failed over, replication stops coming from your production site, since the original environment would be destroyed or inaccessible in a disaster (during such a test, often replication is reversed, therefore, storing changes back in your on-premises environment as if it were the DR target). This is clearly a more extreme type of testing, but it isn't uncommon among companies who want complete assurance that their DR plan will work.

Work with your DRaaS provider to triage any issues that emerged in the test. Often, it's a network configuration or a new system that simply needs updating.

You're not alone – iland is happy to help

As we walk through the steps in building a business continuity plan for your business, you may find you have questions along the way – about technology options, system requirements or the fail-over process itself. Whatever your questions, sometimes phoning an expert can save time and reduce frustration.

“Work with your provider until you're confident your DR solution is working well. Only then, can you declare yourselves prepared for anything nature or technology can throw.”

With over 10 years of experience architecting and delivering Disaster Recovery solutions for our global customers, we've seen it all – from complex networking to legacy systems to challenging procurement processes. In that time, we've learned that the best way to support customers is to roll up our sleeves and help solve their DR challenges from the up-front planning and architecture phase through deployment and even during the disaster itself.

iland has solutions to address your physical and virtual systems, supports co-location and bare metal provisioning, delivers cloud-based backup and an exceptionally secure and well-managed cloud environment, in the event of a disaster. With a slate of industry accolades, partnerships with the leaders in DR technology like Zerto and Veeam, and 24x7 technical support, we're committed to exceeding your DRaaS expectations.

At any point in the process, consider:

Exploring our offerings at <http://www.iland.com/services/cloud-disaster-recovery/>

Read up on why iland was named a Leader by Gartner 2 years in a row in the "[Magic Quadrant for Disaster Recovery as a Service.](#)"

[Speaking with an iland representative](#), who will gladly answer questions, discuss your particular needs, and architect a solution unique to your company

We look forward to hearing from you!

About Veeam Software

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Enterprise™, a business that must operate 24.7.365. To address this, 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 fundamentally new kind of 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 has 53,000+ ProPartners and 282,000+ customers with the highest customer satisfaction scores in the industry. Headquartered in Baar, Switzerland, Veeam has offices in more than 30 countries. To learn more, visit <https://www.veeam.com> or follow Veeam on Twitter @veeam.

Contact

United States	+1 713 868 2267
United Kingdom	+(44) (0) 20 7096 0149
Netherlands	+(31) (0) 10 808 0440
Singapore	+(65) 3158 8438
Australia	+(61) (0) 2 9056 7004



About iland

iland is a global cloud service provider of secure and compliant hosting for infrastructure (IaaS), disaster recovery (DRaaS), and backup as a service (BaaS). They are recognized by industry analysts as a leader in disaster recovery. The award-winning iland Secure Cloud ConsoleSM natively combines deep layered security, predictive analytics, and compliance to deliver unmatched visibility and ease of management for all of iland's cloud services. Headquartered in Houston, Texas and London, UK, iland delivers cloud services from its data centers throughout the Americas, Europe, Australia and Asia. Learn more at iland.com.
