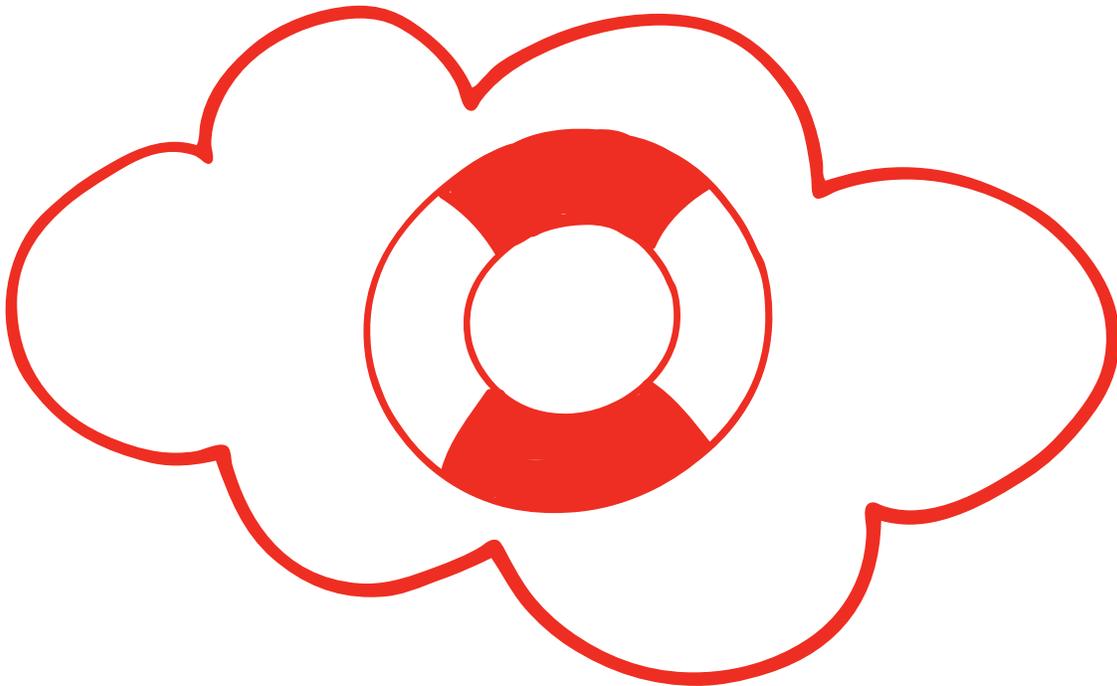


Advanced Technologies Make Thinking About Disaster Recovery a Lot Less Scary.

Why You Should Reevaluate Your DR Strategy Today.



Having survived today, it is human nature to want to have another day just like it tomorrow. Many IT departments, in fact, are set up according to that fond illusion. Having put out the usual fires with the usual heroics, too many IT departments are surprised by the inevitable when a disaster they could have seen coming changes everything.

It could happen. Stories abound about organizations that thought they were prepared only to find out they weren't—a chipmaker that was unable to place orders for two weeks, a hospital that was hacked into by an ex-employee or a retail chain without a functioning point-of-sales system.

Scary and costly stuff. Like a snowball rolling downhill, the cost of downtime increases rapidly over time as the impact spreads across related systems. The bigger you are, the faster it rolls. Research indicates a typical enterprise loses between \$80,000 and \$100,000 for each hour of downtime. So many data disaster statistics have been passed around over the years that the validity of any of them has become suspect. Suffice it to say that in today's online, always-on economy, if your IT systems are down, you're offline—and your competitors aren't.

Floods, fires, hurricanes, a reckless backhoe operator—disasters can be caused by a wide variety of sources. Many disasters are acts of nature, but a surprising number of disasters are caused by people, ranging from sabotage, viruses and incompetence to good old-fashioned, honest mistakes.

Further compounding the problem, mission-critical applications are increasingly dependent on other applications, in a web of interdependence that functions like a living nervous system. You can't let one piece die and expect the rest of the system to survive.

Range of Options

Disaster recovery (DR) typically turns up—or should—during discussions about data center upgrades. Many organizations have put off thinking about updating their data centers for so long that they may not be aware of the wide range of options that are currently available for both production data centers and DR sites.

IT departments that want to keep their DR sites under their control can locate them a safe distance away from the production data center, either in a regional office or other corporate facility or at a co-lo site that they maintain. Because virtualized and consolidated data centers today require so much less real estate than traditional data centers, some organizations have moved the production data center into a co-lo or data center-as-a-service (DCaaS) facility and brought their DR site in house with repurposed gear. There are many choices—and that's without talking about all the cloud options (more on those to come).

For all the sophisticated technologies available today, according to analyst reports, approximately 75 percent of the world's data is still protected by copying it to magnetic tape and shipping it off to some secure offsite storage.

Because it is so pervasive, tape has established a false cost comparison to replication technologies that have emerged in recent years. Tape is relatively cheap until you begin to add in the cost of the time and effort it takes to retrieve your data—recovery point objective (RPO)—and reunite it with the systems that are required to do anything with it—recovery time objective (RTO). It can take hours, if not weeks, to restore data from tapes, and by then, of course, your data is already hours or weeks old.

The cost of downtime per hour across US industries:

Brokerage Service	\$6.48 million
Energy	\$2.8 million
Telecom	\$2 million
Manufacturing	\$1.6 million
Retail	\$1.1 million
Healthcare	\$636,000
Media	\$90,000

Source: Network Computing, the Meta Group and Contingency Planning Research

There is arguably a role for tape backup in your DR strategy, but DR technologies based on data replication, by contrast, provide a continuously updated copy of critical data from one location to another over a storage area network (SAN), LAN or local WAN, so you always have multiple up-to-date copies of your data. Replication often works in combination with data deduplication, virtual servers and cloud computing.

Replicating data between the primary and DR sites before failures occur ensures that data and applications are current at the secondary site. Failover to the DR site happens automatically when the primary site fails. Failback is when the system is restored to its original state before the failure.

Virtual Assist

Virtualization at the server, network and storage layers dramatically facilitates the replication of data within a disaster recovery system. Before virtualization, organizations needed to have the same hardware at their primary and disaster recovery sites. With virtualization, IT can abstract the systems and have different types of hardware at multiple sites. It is also possible to perform tasks such as automating rapid virtual machine rebooting, replicating virtual machines at the hypervisor layer with heterogeneous storage, and turning backups of physical or virtual machines into bootable virtual machines.

There are two types of data replication products: synchronous and asynchronous. Synchronous replication writes data to the primary and secondary sites at the same time. Asynchronous replication involves a delay before the data gets written to the secondary site.

Both types have advantages and disadvantages. While data is always current between sites with synchronous replication, it is more expensive than asynchronous replication, introduces latency that slows down the primary application, and only works over distances of 30 miles to less than 200 miles. Synchronous replication is preferred for applications with low RTOs.

Asynchronous replication is designed to work over greater distances. It requires less bandwidth, and it is often a better option for disaster recovery. The data at the target device is never exactly up-to-date with the source data, but the difference is typically a matter of minutes. Replication tools such as Double-Take by Vision Solutions and Site Recovery Manager (SRM) by VMware are asynchronous.

The trend toward converged infrastructures, coupled with advances in software defined storage (SDS) and software defined data centers (SDDC), promise to further facilitate DR by increasingly liberating not only data but whole platforms from underlying hardware. Cloud computing technologies such as automation and orchestration will also enable dynamic environments that go a long way toward taking care of themselves. Cloud computing even liberates IT departments from having to buy, install, maintain and support their own hardware.

There are two types of data replication products: synchronous and asynchronous.

At its most basic level, choosing the right disaster recovery technology is a business decision.

DR in the Cloud

Several different models for cloud-based DR are now available, including:

- *Do-it-yourself cloud DR.* For those IT professionals who are very confident in their expertise in both cloud and DR, the “do-it-yourself” approach is a flexible and cost-effective way to take advantage of public cloud resources for DR purposes. Keep in mind that cloud disaster recovery services involve more than just cloud storage; they encompass planning, process, integration, testing and constant vigilance.
- *DR-as-a-Service (DRaaS).* DRaaS solutions are prepackaged services that provide a standard DR failover to a cloud environment. Customers can buy these services on a pay-per-use basis with varying rates based on required RPOs and RTOs. Service providers either deploy agents to replicate data and applications or use image-based backups to send data to the cloud.
- *Cloud-to-cloud disaster recovery (C2C DR).* C2C DR, the least common of the cloud scenarios, is the ability to failover infrastructures from one cloud data center to another, either within a single vendor’s environment or across multiple vendors.

Cloud-based DR is available in three different architectures:

- *Cold cloud recovery* is a recovery cloud environment that contains backup images of the environment that must be first rehydrated before recovery. In this scenario, the cloud resource that is consumed (meaning you pay for it continuously) is storage. Achievable RTOs are usually six to 48 hours and RPOs are 24 to 48 hours.
- *Warm cloud recovery* is a cloud that contains up-to-date versions of production virtual machines (VMs) that are kept offline and idle. During a disaster or test, you can quickly spin up VMs from offline VMs, resulting in RTOs usually of two to six hours and RPOs of minutes to a few hours.
- *Hot cloud recovery* is an alternate site running a replica of the primary site that is not in use unless the primary site fails. With a hot site in the cloud, you would constantly run VMs dedicated to DR and keep them up-to-date using replication or backups, depending on the desired RPOs. The RTOs in this approach are usually less than one hour.

A Business Decision

Each of the DR replication technologies has its own set of pros and cons. Deciding which technology—or combination of them—is right for your organization, however, involves more than just technical considerations.

At its most basic level, choosing the right disaster recovery technology is a business decision. As a result, before IT departments can decide on a disaster recovery technology, they have to establish the business priorities and objectives that the technology will be required to meet. This step can be one of the most challenging for IT because it involves reaching across the aisle to the business side.

“Immediately” and “none” are the two most common answers from department heads when asked how fast data needs to be recovered (RTO) and how much data they can afford to lose (RPO).

Establishing appropriate, realistic and affordable RTO and RPO levels needs to be negotiated between the IT department and business leaders. It's often a balancing act. Having an outside, unbiased business continuity planner present during these discussions helps them proceed to a mutually acceptable conclusion.

Once realistic recovery objectives are identified, developing a tiered strategy that meets the specific requirements and budgetary constraints of your organization is relatively straightforward. Options abound for tier-one mission-critical applications that need some form of automated failover (short RTO). Less expensive technologies are available for applications that can be recovered more slowly.

To clearly establish the "You are here" point of reference, a disaster recovery team needs to include technology experts and business analysts who conduct a detailed evaluation of your entire IT infrastructure, as well as the people, data, processes and technology that your organization depends on to function effectively.

The two key components of a DR assessment are:

- A *Risk Assessment* that involves identifying a potential risk event, assessing the likelihood of the event occurring, and defining the severity of the event's consequences. Risks could be anything from a power outage or hardware failure to a tornado or flood.
- A *Business Impact Analysis* that evaluates mission-critical business functions and identifies and quantifies the impact a loss of those functions (e.g., operational, financial) would have on the organization.

Not Optional

Even if you haven't been thinking about all the various ways your IT could suffer a significant outage, a growing list of regulatory agencies have been thinking about it for you, and they are drafting regulations requiring you to be able to demonstrate that you are prepared to respond to, if not prevent, disaster. PCI DSS, HIPAA, FISMA and ITAR all require a functional—and routinely tested—disaster recovery plan. If industry regulators can't make you think about DR, a growing number of companies that depend on partners for supplies are now asking their suppliers to prove they are prepared for disaster—or risk terminating their relationship.

Pushed toward a viable DR strategy by regulators, partners and their own long-term best interest, many IT departments, at the same time, are held back by a data center full of aging technology and a shrinking capital budget. DRaaS offers a way out of this dilemma by sidestepping the investment in new technology and shifting DR expenses from capital budgets to operational budgets. For example:

Total Safety of Houston is a leading global provider of integrated industrial safety services, strategies and equipment. An organization with an intense focus on safety, Total Safety was acutely aware of the risks it faced by having its data center in "hurricane alley."

A disaster recovery assessment had shown that the only way to provide a totally safe data center environment would be to build a duplicate disaster recovery site outside Houston. Total Safety Technical Operations Manager Kevin Croteau had started investigating options for a remote site, but he soon found every offsite option becoming cost prohibitive.

Key Business Benefits of DRaaS

- Ability to shift costs from a capital expense to an operational expense model
- Flexibility to expand or reduce service levels as business needs change
- Access for your IT team to specialized and highly skilled personnel as needed
- Freedom for your IT team to concentrate on dealing with the physical and personal challenges presented by the disaster while the DRaaS provider spins up critical business systems
- Peace of mind

The DRaaS Development Process

Analysis and Design

- Analysis of server requirements
- Analysis of network and user requirements
- Design cloud target environment and network access
- Design data migration plan
- Design test criteria

Implementation

- Implement cloud target infrastructure and install software
- Install replication software
- Review and implement parameters
- Backup data to mobile array for initial seeding
- Discuss application-related data replication sets and connections
- Discuss system-related replication sets and connections
- Discuss initial mirrors and difference mirroring
- Connect replication sets, begin mirroring and monitor replication-related activity
- Verify software is operational on source(s) and target(s)

DRaaS offered Total Safety an affordable solution. Instead of investing in a DR site and a duplicate set of servers and storage, Total Safety was able to implement a disaster recovery site within a warm cloud environment at a Logicalis enterprise cloud facility in Cincinnati.

Deploying its DR site in the cloud eliminated the risk of downtime and enabled the entire cost of the disaster recovery platform to be paid from Total Safety's operational expense budget, instead of its capital expense budget.

Because Logicalis already had the hardware and software systems in place, implementing this solution took a fraction of the time of building a comparable disaster recovery site.

Total Safety's DRaaS solution is basically a private cloud, disaster recovery and failover site that is part of its MPLS network. A VPN/Internet connection provides a secondary route, eliminating a single point of failure.

The solution was designed to minimize monthly expenses by leveraging replication software. Until there is a disaster, Total Safety only pays for the storage and compute capacity required to replicate the data—which amounts to less than half of what it would cost if their whole data center was running on the Logicalis cloud platform.

In the event of a disaster, VMs can be spun up in a matter of minutes to keep all Total Safety systems running smoothly in the cloud, allowing Total Safety IT staff in Houston to focus on the situation on the ground at home.

DRaaS Design and Onboarding

Designing a DRaaS solution needs to be a collaborative effort between the client's IT department and the service provider. IT departments taking advantage of DRaaS are not excused from providing a great deal of information during both the design and onboarding stages.

Documentation must identify the operating systems, who will recover the data, where it will be recovered, how connections will be made, how the failover will be tested, how the data will be verified, and the criteria for releasing data to the production environment. The key players and their full contact information and anticipated response times must also be determined.

Even after the matching infrastructure is set up in the cloud, it is still the responsibility of the client organization to validate that the applications in the cloud site are fully functional.

DRaaS solutions need to be designed to adhere to RPOs and RTOs based on business considerations and well-defined processes for recognizing and declaring a disaster. They also require tight integration among the network, firewalls and load balancers, and also the Web, database and storage tiers. If DRaaS is not equipped to recover the entire stack above storage, the recovery will fail.

A critical determination in the design of a DRaaS solution is an accurate understanding of capacity requirements. Storage capacity is straightforward: You need the same amount of storage in the cloud as you have in the data center being replicated.

The trickier calculation is the processing and memory requirements that will need to be met in the event of a disaster when the necessary virtual machines are spun up to run the applications. Keeping a replicated site in readiness typically takes about 15 to 25 percent of the compute power and memory capacity that will be required during a recovery.

Onboarding

During onboarding, the client syncs the systems to be protected to a mobile array that is delivered to the target site where the data is used to seed the target cloud environment. Once that data is in place, replication only has to copy the deltas (changes) over the WAN. Besides new data, all upgrades and patches to OSs and applications at the source data center are replicated to the target DR site in the cloud.

Different replication technology is required for different IT environments. With the right tools, it is possible to replicate both physical and virtual environments and leverage virtual replication technologies for virtual environments.

Beyond technology, designing a DRaaS solution also requires an understanding of how users access and use their applications. A technically successful recovery will still fail if users get their data back in a timely fashion, but not in any form they can use to do their jobs. Networking configurations have to be set up to ensure that users can access their applications during a recovery.

Social media and the “bring your own device” (BYOD) trend present both challenges and solutions in the event of a disaster. On the plus side, users today have a wide variety of mobile devices they can use to access critical systems during a disaster. On the negative side, critical data is increasingly distributed, if not scattered, throughout an organization’s extended IT infrastructure. Without a well thought out BYOD strategy, it doesn’t take a natural disaster to suffer a major data loss. At the same time, although Facebook, Twitter and other social media platforms make it possible for employees to communicate with each other, the kind of open-ended communications those platforms enable could also make the disaster worse by contributing to widespread fear, uncertainty and doubt (FUD).

Cloud to Cloud

Cloud-to-cloud data protection services take the DRaaS model to the next level by allowing organizations to run both their production data center and their DR site in the cloud and shift the total cost of running and protecting their critical systems from capex to opex. No more having to buy new technology only to have to update and refresh it every few years. C2C is often bundled with monitoring and management, so the IT department is liberated from mundane operations as well.

A healthcare ISV, for example, hosts applications for its clients in the Logicalis Enterprise Cloud in Cincinnati (LEC East). In the event of a disaster in Cincinnati, its applications and associated data automatically failover to the Logicalis Enterprise Cloud in Phoenix (LEC West). Similarly, the e-commerce sites that a national retailer runs in LEC East are set up to failover to LEC West.

It is entirely possible to mix and match which systems you want to support in your own data center and which ones you run and/or failover to a cloud environment. In each of the above examples, the cloud clients maintain development environments outside of the Logicalis cloud.

What is, or should be, included in a DRaaS solution?

- A multi-tenancy cloud environment in a Tier 1 data center
- Geographic diversity and data center redundancy
- A security framework based on industry best practices, i.e., HIPAA, PCI DSS, FISMA, ITAR and MASS LAW compliance
- SLAs with true financial penalties
- Dedicated disaster recovery delivery managers
- Guaranteed reservations for additional capacity required in the event of a disaster
- Monitoring and management based on ITIL best practices
- Quarterly review and an annual full recovery test

If you don't
already know,
you should
consider a DR
assessment...
before it's too
late.

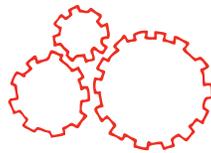
New Paradigm. Old Concept.

It was one thing to leave IT systems exposed to the risk of a company-destroying data disaster when preventing disaster cost only slightly less than having one. Today, some combination of virtualization, replication, co-location, DCaaS, DRaaS and C2C make at least some level of data protection a no-brainer.

What is too often lacking is the willingness of IT departments and business managers to face the potential for data disaster in their organizations and make building a strategy to prevent it a priority for which they are willing to pay. Unfortunately, not thinking about disaster will not prevent it from happening, nor will it help you recover from it.

At its most basic level, a disaster recovery strategy is not about technology—or, at least, not just about technology. It's also about people and processes and the confidence to look into the future without turning away from the potential for bad things to happen.

The good news is that with the new generation of disaster recovery technologies, it is possible to confidently prepare for disaster recovery in a way and at a cost that can be tailored to your organization's tolerance for risk. You just have to decide how much risk you are prepared to take. If you don't already know, you should consider a DR assessment...before it's too late.



 **LOGICALIS**
Business and technology working as one

866.456.4422 | www.us.logicalis.com