How to Improve Business Resiliency with Disaster Recovery-as-a-Service (DRaaS)

The fundamental premise behind the very existence of an IT department is the expectation that systems will be kept operational and business-critical applications will always be available. With that in mind, IT leaders invest an immense amount of time and capital into ensuring that their systems are as robust as possible, data is protected and performance is monitored.

However, as the pace of business and change accelerates, it is an increasingly difficult task to stay ahead of the curve. This is especially true when trying to plan for disasters, which are, almost by definition, unforeseeable. Disasters can come in many forms—from a site-wide outage due to a natural event to data corruption caused by human error to malicious attacks requiring extraordinary measures to overcome and recover to a safe point—and the number and types of events that can cause disruption are on the rise.

- More than 50% of organizations declared a disaster and failed over operations to their recovery site at least once over the last five years.[1]
- In 2020, ransomware incidents will grow as attackers learn that holding data hostage is a quick path to monetization.[2]
- In Q4 2019, the average downtime caused by a ransomware attack was 12.1 days—nearly double what it was at the same time in 2018.[3]

What's more, the stakes have never been higher.

The potential implications of disaster-related IT failures are far-reaching and can range from breached medical records to legal repercussions, privacy leaks, lost revenue and damage to brand value. For all these reasons, any business reliant on IT for any part of its ongoing operations needs to have an up-to-date, regularly tested disaster recovery plan.

When considered holistically, a disaster recovery plan includes both near real-time replication capabilities as well as prevention and resiliency strategies. Well-planned and properly executed, a comprehensive disaster recovery plan forms the cornerstone for maintaining business operations in the face of any crisis. Unfortunately, disaster recovery plans too often fall short in terms of scope or practical application. There can be many reasons for this, but the most commonly cited are cost and lack of available internal resources.

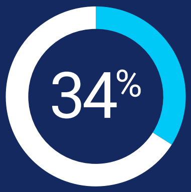

Only 34% of organizations say they are somewhat prepared to recover their data centers in the event of a site failure or disaster event.[4]

Many businesses view a disaster recovery plan as an insurance policy against an unlikely event and fail to adequately protect their assets. Companies often assume that their legacy policy of either backing everything up to tape and shipping it offsite or remotely replicating it to a storage archive is sufficient. This kind of simple backup and recovery strategy forms the backbone of many disaster recovery plans; however, a robust disaster recovery strategy that can ensure continuity of business operations requires a careful examination of whether the recovery time and recovery point capabilities of traditional backups are sufficient for the needs of that business. In addition, many in-house backup and recovery solutions cannot be geographically dispersed effectively, nor can they efficiently take advantage of real-time data replication technologies. The cost of real estate—let alone the equipment, staff and operational overhead of a secondary site—is also prohibitive. Even in the case of existing secondary office locations, the expense to build out and manage a second data center is simply not practical.

DRaaS can replace or augment traditional disaster recovery planning

This financial catch-22 of downside risk versus the high cost of up-front capital investment has led to the rise of disaster recovery-as-a-service, also know as DRaaS. This service leverages the broad reach and cost-efficient, on-demand capacity of the cloud along with partner-provided disaster recovery expertise. Additionally, when sourced as a service, disaster recovery becomes a manageable operational expense rather than what is often a substantial upfront capital expenditure that requires significant, short-term budget trade-offs.

Similar to selecting a cloud service provider, deploying an effective DRaaS solution should begin with a solid understanding of which business and technical goals are driving the need for a recovery plan. This means starting with a business impact study conducted in house by an auditing group or, if offered, by the chosen DRaaS provider itself. This research should be undertaken in cooperation with the business leaders or business units themselves—not just the various IT departments—to classify the processes and systems most vital to the success of the company. These classifications will be based on agreement regarding the acceptable amount of downtime and data loss recovery period for these systems. The goal is to understand not only the impact to the business processes, but the expectations of the people responsible for each distinct aspect of the identified business operations.

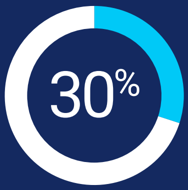

Only 30% of applications and data fall into the mission-critical category and require low RTO and RPO strategies.[5]

It is important during this assessment to ensure that all parties clearly understand the difference between the concepts of recovery time objectives (RTO) and recovery point objectives (RPO) as well.

RTO and RPO: Understanding the value of time and data

> **RTO:** The time that an application or business process is unavailable. For revenue-generating activities, this is easily measured in terms of lost income such as sales per hour of downtime. For operational activities, such as manufacturing production, the measurement is often calculated in terms of lost productivity, the write-off of perishable product or SLA penalties.

> **RPO:** How much data is lost after implementing recovery from a backup or during failover to a secondary site. This is measured most simply in terms of the value of any lost data, such as information collected about purchasing, inventory or lost sales. There are often secondary costs to data loss as well, including compliance penalties, public relations expenses and loss of brand value, that can easily exceed the value of the data itself.

At first pass, many business leaders will focus most of their energy on the amount of time a given system, process or application can be offline before being recovered, often insisting that even a minor disruption represents a catastrophic loss of revenue or productivity. This downtime window is referred to as the recovery time objective and generates a lot of interest because it is the impact most easily perceived by people trying to get work done. Days, hours, minutes or even seconds of downtime will translate into lost revenue, contractual penalties or loss of return on investment due to missed milestones in key projects. Consider a business reliant on online or mobile sales transactions such as a large retail enterprise. For a company like this, even a single minute of its sales site being offline translates into hundreds of thousands or even millions of dollars in lost revenue.

Recovery time can inconvenience a business and even cause significant financial impact, but **poorly planned recovery points can completely cripple a company.** In addition to lost revenue and decreased productivity, lost data can also create financial or legal liability, regulatory compliance penalties and diminished confidence in the business.

Systems with a low tolerance for downtime need to be classified as such, with effective process recovery plans in place. These plans can include straightforward solutions such as the ability to reroute workstreams to warm or hot disaster recovery sites or implementing load-balanced, geographically distributed access to ensure that workers and customers are not left unable to transact business. This is important not only in the case of a disaster or malicious acts but also for more mundane things like system maintenance, software patches and infrastructure updates to ensure that this routine, behind-the-scenes work is minimally disruptive to the flow of business operations.

However, while the RTO is very important, the recovery point objective cannot be overlooked or underestimated. It is here that many businesses can be truly lost as they deal with far-reaching implications not initially understood at the time of a system outage. The RPO defines the amount of data that can be sacrificed in order to secure a clean recovery of business operations. A daily backup scheme, for example, might mean that up to 24 hours of data could be entirely lost. This affects everything from supply chain tracking to order and manufacturing processes to financial operations. It is not difficult to imagine the impact of a day's worth of customer credit card transactions lost within an online ordering system that had to be restored to a recovery point a day old. These transactions would likely have already been processed against customers' banks, but the business would be left without vital information like the customer's name, what they ordered and where the item should be delivered.

Recovery time can inconvenience a business and even cause significant financial impact, but poorly planned recovery points can completely cripple a company. In addition to lost revenue and decreased productivity, lost data can also create financial or legal liability, regulatory compliance penalties and diminished confidence in the business. Ensuring that business leadership not only understands the terminology surrounding RTO and RPO, but also completely thinks through the business implications of each, is a critical first step in developing an impact study and classifying each workstream.

Additionally, the applications and processes identified as business or mission-critical during the initial assessment need to be considered in relation to each other and dependencies on other IT systems. Many core applications are heavily reliant on a variety of systems, tools and even legacy applications that have not been effectively modernized. These ancillary systems are used for a variety of easily overlooked, but nonetheless critical, processes such as login and user authentication, data storage and retrieval and network access. Effective recovery from a disaster means considering all of these dependencies and implementing a plan that is broad-reaching enough to cover them all, yet flexible enough to adapt as the flow of information changes over time.

Think again of that online consumer transaction. An order on the website flows from a web front-end through a customer relationship management (CRM) database and transaction processing application to an order management utility. That, in turn, is subsequently tied into a supply chain or inventory control system, possibly into a manufacturing resource manager and eventually into the back-end accounting and finance systems. At each step of this transaction, the various applications that move an order through the business have infrastructure dependencies, access controls and data collection and archiving processes that are tied together. These cross-dependencies need to be fully explored and categorized so that keystone systems and infrastructure are not categorized lower than the more obvious, mission-critical applications they support. Thus, it is important to think in terms of the entire business process, not just single systems or applications. A traditional backup application is typically focused on a specific data set tied to a specific application and may not take into account the spectrum of different systems relying on that data. Application failover or clustering functions may help ensure that a core application stays up and accessible but may break key upstream or downstream dependencies in the event that failover occurs, unless scrupulously planned and implemented.

DRaaS can help safeguard the entire workstream from start to finish. Unlike traditional backup utilities, a disaster recovery service is built around an entire business process that takes into account cross-dependencies, performance requirements and the implications of both downtime and recovery points.

The best disaster recovery scenario avoids having to recover anything at all

For all this planning, avoiding an outage in the first place should be the primary goal of any DR strategy. This doesn't mean avoiding data protection, backups and recovery mechanisms, but instead building resilience into the systems and operations in such a way that recovery from backup is truly a last resort. By implementing disaster recovery as both a plan and an ongoing service, enhanced protections can be applied up front to help mitigate and manage disasters before they require system downtime or data recovery.

This requires implementing intrusion detection policies that are consistently monitored and upgraded to handle ever-evolving threats. Malware protection and isolation options are also required, as are distributed workloads with real-time replication such as a DRaaS solution. Finally, there must be regular testing, auditing and updating of the DR plan as the needs of the business change, applications grow or are updated, new technologies are adopted and as new threats and malicious actors appear on the scene. It is a continuous cycle of planning, monitoring and adapting in an environment where the pace of change only ever seems to accelerate.

Beyond these preventative measures, regular review and testing of a disaster recovery plan ensures that critical exposures or new dependencies are not uncovered during an actual incident. An advantage of DRaaS is the ability to conduct tests in a controlled manner. With the right partner, change management can also be planned and coordinated as new systems are brought online, new processes are established or legacy applications are retired.

DRaaS offerings go well beyond what most companies are capable of delivering in house. They bring to bear extensive expertise via certified consultants and operators who can deliver a bespoke approach to tailoring DRaaS services to the unique needs of each of their clients. This avoids the traps of one-size-fits-most or out-of-the-box offerings, while still allowing for a wide range of simple to complex services to effectively implement disaster recovery for any size business.

DRaaS offerings are not a one-size-fits-all service. DRaaS offerings bring to bear extensive expertise via certified consultants and operators who can deliver a bespoke approach to tailoring DRaaS services to the unique needs of each of their clients.

While DRaaS offerings have proliferated in recent years, many of these programs are only self-service and rely on their customers having the in-house knowledge and expertise to select the right options and services with minimal assistance. Often, low-cost DRaaS providers offer little or no customization of their services, instead, relying on predefined, generically sized packaged offerings that may not be a good fit for many businesses. These services may also fail to fully account for systems and processes that are not natively hosted in the cloud. A colocation provider allows not only for cloud-based services but also hybrid solutions and single tenant hosted services—even for applications or services that cannot be easily virtualized away into the cloud or require unique physical components like phone systems.

The challenges of implementing a DRaaS strategy can, at times, seem overwhelming. Often the biggest obstacle is simply making the decision to prioritize it in the face of competing business needs and agendas. **However, the face of disaster is changing, and disaster events are getting increasingly costly every year. The stark truth is that waiting or delaying disaster recovery planning can have consequences that are all too often more than a business can withstand.**

[1] The State of Disaster Recovery Preparedness 2020, Forrester Research; [2] 2020 Predictions, Forrester Research; [3] Cybersecurity Ventures; [4, 5] Survey: Forrester/DRJ State of DR Preparedness

About Flexential

Flexential empowers the IT journey of the nation's most complex businesses by offering flexible and tailored solutions in colocation, cloud, data protection, managed and professional services. The company builds on a platform of three million square feet of data center space, in 20 highly connected markets, and on the FlexAnywhere^{TM} 100GB private backbone, to meet the most stringent challenges in security, compliance and resiliency.

Visit www.flexential.com.

Follow Flexential on LinkedIn, Twitter and Facebook.