RICOH
imagine. change.

Business Continuity in the Cloud Era

How Disaster Recovery as a Service (DRaaS) is Giving Businesses a Critical Edge

Introduction: Disaster Recovery Moves Into the Spotlight

In early 2020, the coronavirus pandemic rapidly disrupted every sector of the business world in broad and complex ways. The sudden shift to remote work and decentralized collaboration disrupted employees' access to critical networks, systems and applications, threatening productivity and impacting companies' abilities to connect with and serve their customers. This disaster was a lot different than what most companies plan for — technology failures, natural disasters, ransomware, etc. Nevertheless, the pandemic undoubtedly brought the importance of Disaster Recovery (DR) and High-Availability (HA) directly into the spotlight as never before, making 2020 a turning point for DR and HA planning. The ability to maintain continuous digital operations is increasingly a critical competitive differentiator for businesses in every segment — not just enabling a business to survive, but enabling a business to thrive by guaranteeing uptime, maximizing performance and earning priceless customer loyalty by delivering consistent, reliable service.

This report will look at six emerging insights that shed light on this new paradigm of business continuity. We'll look at what's changed — from market expectations, to what businesses are protecting, to the most common business continuity risks. And we'll look at how new cloud-powered technologies and as-a-Service solutions are making higher standards in DR and HA accessible, practical and achievable for more businesses.

Six Emerging Insights on Modern Business Continuity Planning

1. Disaster recovery planning is more important than ever before
2. Everyday disasters are the biggest risk to continuity in business operations
3. What's "business critical" versus "mission critical" has changed
4. Business continuity planning processes are the same, but the goals are higher
5. Cloud-based disaster recovery technologies are driving improved recovery objectives
6. The "as a Service" delivery model is unlocking flexible, practical access to next-generation DR solutions

Insight #1: Disaster Recovery is More Important Than Ever

The concept of disaster recovery goes all the way back to ancient Babylon, some 6,000 years ago, with the issuing of merchant insurance policies that covered lost goods in the event a ship sank. That is to say, businesses have always recognized that having a disaster recovery (DR) program in place provided essential protection — against lost revenue, lost customer trust, reputation damage, etc. But in the modern business world, DR is about much more than recovering the cost of lost goods in the event of a disaster — it's about rapidly recovering the ability to serve your customers. That's because, in our knowledge economy, the primary "goods" produced by most companies are expertise and trusted service. And there is no recovering these intangible goods; today, downtime increasingly means permanent losses.

Expectations for Business Continuity Keep Rising

Not only are the permanent costs of downtime increasing, the expectations for businesses to deliver continuous, consistent service are higher than ever. We all live in a digital world of instant gratification, and this culture carries over to customers in every segment who have right-now expectations of 24/7/365 service. Customers' have rapidly decreasing willingness to tolerate downtime and service interruptions — and falling short of their expectations has a more immediate and damaging effect on trust, loyalty and overall brand reputation. But the DR pressure doesn't just come from customers — increasing regulatory requirements around customer data protection, data privacy and data recovery make high-performance DR essential to compliance, as well.

The Digital Ecosystem is Growing More Complex

Perhaps most challenging, these rising standards intersect with the rapid expansion of the typical business' digital ecosystem. Companies are managing vastly more applications, much broader networks and exponentially more data than even just five years ago. Moreover, this digital ecosystem is expanding to include more public-, private- and hybrid cloud-hosted applications and systems, making it more challenging to build comprehensive DR programs that manage and protect all these cloud-hosted applications and systems.

Business Continuity Glossary

- **Business Continuity:** Business continuity describes the processes and procedures an organization puts in place to ensure that essential functions can continue during and after a disaster. Business continuity planning seeks to prevent interruption of mission-critical services, and to reestablish full functioning as swiftly and smoothly as possible. Business continuity strategies typically include two main components: High Availability and Disaster Recovery.

- **High Availability (HA):** High Availability, or HA, refers to a system or component that is continuously operational for a desirably long length of time. Availability can be measured relative to "100% operational" or "never failing." A widely-held but difficult-to-achieve standard of availability for a system or product is known as "five 9s" (99.999 percent availability).

- **Disaster Recovery (DR):** Disaster Recovery, or DR, focuses on duplicating/recovering computer operations after a catastrophe occurs, such as a fire or earthquake. It includes routine off-site backup as well as a procedure for activating vital information systems in a new location.

The Harsh Realities of Disaster Recovery

70.9%
of companies have experienced **unplanned downtime** in the last 2 months[1] with an average cost of
$740,000

96%
of companies have experienced **at least one system outage** in the last 3 years[2]

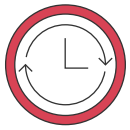

One Hour of Downtime costs ....

$25,000
for a small company[3]

$540,000
for a large enterprise[4]

which equates to **$9,000 per minute** for larger enterprises[4]

3 in 5
businesses are **not confident** they can meet Service Level Objectives to fully recover systems and data

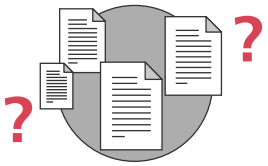

84%
of businesses say their current data protection solution **wont meet future business challenges**

Remote Work Amplifies Need for Cloud-App Uptime

The coronavirus pandemic drove a categorical shift to remote work and decentralized collaboration for businesses in every segment. Even as the acute risks of the pandemic ebb, this shift to remote, flexible work will remain in many companies — making businesses fundamentally reliant on their remote employees' continuous access to cloud-hosted applications. To fully take advantage of the competitive advantages that can be gleaned from enabling an agile, flexible, largely remote workforce, businesses need to intensify their focus on protecting and guaranteeing uptime of these key hosted applications.

Insight #2: Everyday Disasters are the Biggest Risk

When you consider disasters that could put your business continuity at risk, it's common to immediately think of the kinds of once-in-a-generation disasters that have brought entire economies to a halt: events like 9/11, Hurricane Sandy ravaging the East Coast, and now, the coronavirus pandemic. But the reality is that the biggest risks to most businesses are much more mundane. The biggest culprits of downtime are human error, security breaches, hardware and software failures.[6] Simple user error accounts for an enormous number of downtime incidents every single day. And the risks from user error are only growing as modern technologies empower employees with greater access to data, networks and systems as they work, share and collaborate in new ways. Along with user error comes what is now the biggest source of downtime: ransomware — where an external actor finds a way in (often through a phishing or 'click the link' attack) and holds a database/system hostage for ransom. Or, to think of the relative risks another way, isolated, acute power outages add up to millions more in lost revenue than large-scale outages like those from a hurricane.

The Pareto Principle: Following the 80/20 Rule in Disaster Recovery Planning

Anyone in the business world has undoubtedly heard of the 80/20 rule in some context. Also called the Pareto Principle after the Italian economist who first coined the idea, the 80/20 rule stipulates that 80% of consequences come from 20% of the causes. Just as businesses use the 80/20 rule to guide sales, marketing and operational strategies, so, too, should they leverage the 80/20 rule to guide their disaster recovery planning: Focus on preparing for the small set of common, relatively mundane occurrences that cause the vast majority of downtime incidents and damages.

Insight #3: What's "Mission Critical" has Changed

It's fitting that Pareto's 80/20 rule not only applies to what you're protecting your business from, but what you're actually protecting. In the typical organization, roughly 20% of your business applications account for 80% of your business value. These are your "mission critical" applications — the core of your disaster recovery and high-availability focus.

Mission-Critical Capabilities in the Modern Enterprise

A decade ago, financial applications accounted for the critical 20% in most businesses. But today, the typical organization's mission-critical areas have changed:

1. **Communication & Collaboration Platforms**

   Modern businesses rely on email, mobile applications, remote access and collaboration tools like Google Drive, Dropbox, Slack and more to power their everyday productivity. Disruption of these tools can bring basic internal workflows to a halt — and quickly impact a company's ability to serve its customers.

2. **Revenue Generating Systems**

   The continued digital transformation of the business world has changed where businesses transact with their customers. Today, eCommerce platforms, cloud-aggregation tools, order entry and payment processing systems are how businesses actually deliver services and take payment from customers.

3. **Backend Operations**

   Financial systems remain mission-critical, but they've been joined by ERP systems and other tools that power the backend operations of a modern digital business. Any disruption in these backend systems can disrupt the "digital supply chain" of services flowing internally within the organization.

The Goal of Disaster Recovery: Protect the 20% against the 20%

Looking at the most common sources of downtime issues, and the most critical applications to protect, a simple principle emerges: Protect the 20% against the 20%. That is, protect the 20% of your most mission-critical applications against the 20% of the most common risks.

Mission-Critical vs. Business Critical

These two terms tend to get used interchangeably. But, in fact, there is a distinct — and important — difference between mission-critical applications and business-critical applications:

Mission-Critical

- Effects the entire business
- Outage would rapidly stop all business
- Example: financial and billing systems which facilitate millions of transactions per minute

Business-Critical

- Effects a particular line of business
- Overall business can continue operating during outage
- Example: HR payroll system

Distinguishing between applications which are truly mission-critical and those that are business-critical is an essential part of disaster recovery planning. You want to invest in making sure your mission-critical applications have near-real-time RPO and the shortest possible RTO to minimize the risk of a total interruption in business.

Insight #4: The Process is the Same — But the Goals are Higher

The basic process of building and executing a business continuity program remains exactly the same as it was five, 10 or even 20 years ago. The basic objectives also haven't changed:

The Process

1. **Understand your business:** Start with a comprehensive business impact and risk assessment to put an objective price on discrete threats to availability — and quantifies the value of recovery.
2. **Develop a continuity strategy:** Design targeted, cost-effective responses to the 20% of threats most likely to occur. This includes establishing realistic service-level agreements (SLAs) for your customers.
3. **Implement the strategy:** Deploy your targeted disaster recovery strategies to protect and achieve the defined SLAs.
4. **Exercise and maintain the plan:** Regularly and thoroughly test and review your disaster recovery plan, updating to address new business-critical applications, emerging risks and evolving customer needs.

The Objectives

- **Recovery Time Objective (RTO):** The period of time within which technical services and/or business functions must be recovered and available after an outage (e.g. one business day); measured from the time of disaster to the resumption of production operations.
- **Recovery Point Objective (RPO):** The acceptable level of data loss exposure following an unplanned event. This is the point in time (prior to the disaster) to which lost data can be restored; typically the last backup taken offsite.
- **Budget:** The acceptable level of money that an organization is willing to invest to optimize RTO and RPO based upon their business recovery objectives.

Businesses Targeting Higher DR Performance — at Lower Cost

Companies are still using the same basic process and aiming for the same core disaster recovery objectives. So, what has changed in disaster recovery? Well, five or 10 years ago, most companies built their disaster recovery program around one of the objectives (RTO, RPO, Budget). Today, they're aiming for all three.

Before, some organizations designated service availability as their top goal, focusing on a short RTO to restore service as rapidly as possible. Others recognized that losing data was more damaging than service downtime, so they focused on RPO to ensure they could recover from an incident without losing any of this valuable data. For many organizations, budget became the de facto key objective, as they worked to achieve reasonable RTO and RPO within a set cost.

Now, as business continuity becomes essential not just to survival, but to gaining a competitive edge, businesses recognize that they can no longer afford to focus on just one goal. Instead, they're targeting the intersection point of all three objectives. They're aiming to achieve a shorter RTO to recover rapidly, while trying to enable near-real-time RPO to mitigate data loss. And while most companies are seeing disaster recovery budgets increase, the reality is that the growing complexity of the challenge makes cost-effectiveness as relevant as ever.

Another major change is that the cost to implement strong disaster recovery programs has continually decreased over the last 5-10 years. Deploying a best-in-class disaster recovery and high-availability strategy is vastly more economical than a decade ago. New cloud-based warm- and hot-site approaches can cost as little as 25% of the cost of the production environment. And this complete reframing of the budget side of the equation is allowing companies to effectively target both shorter RTO and near-real-time RPO as reasonable, simultaneous goals.

Insight #5: Cloud-based DR Technologies Driving Improved RTO + RPO

To meet loftier disaster recovery and business continuity goals, businesses know they need more powerful technologies and more sophisticated solutions. Fortunately — as with just about every other aspect of the modern digital business — cloud-based technologies are driving dramatic improvements in the performance capabilities of disaster recovery technologies at all levels. From disk replication, to bandwidth availability, to high-availability solutions, to cloud aggregation, to software defined networks, advanced cloud-driven solutions have changed what's possible in disaster recovery, unlocking radical improvements in RTO and RPO.

Virtual Servers Expand Recovery Options

Increasing adoption of virtual servers, including cloud federation, has opened up a wider range of disaster recovery options and alternatives. This enables organizations to better customize their disaster recovery programs to meet their specific needs and goals.

Hypervisors Drive Cross-Platform Interoperability

Hypervisors from all providers are built with resiliency in mind and are running toward cross-platform interoperability. Interoperable hypervisors are expanding hardware and software capabilities, improving security, reliability and device independence and allowing for complex applications — all while making it easier to administer, manage and control costs on the back end.

Cloud Aggregation and Cross Cloud Orchestration will lead to operational efficiency during a disaster.

Software-Defined Networks (SDNs) Enable Decentralized Systems

The advance and adoption of SDNs enables the extension of a network framework across decentralized physical architecture and multiple locations. This allows disaster recovery programs to be implemented across decentralized work environments, empowering and protecting remote work.

Active Backup Technology Takes Major Steps Forward

Cloud-driven Active Backup technology has led to exponential improvements in Active Backup performance. Cloud-based backup management software has vastly improved the reliability of backup results, while also allowing more centralized, effective backup administration.

Asynchronous Replication Shortens RTO

The computing power and unlimited storage capabilities of cloud-based DR solutions enables asynchronous replication. This approach uses a disk-to-disk copy and maintains a replica of the database or file system by applying changes to the replicating server at the same time changes are applied to the protected server. With synchronous mirroring, the RTO can be minutes — ideal for critical applications that can accept little or no downtime or no data loss.

Multisite, Multi-cloud Replication Optimizes RTO

Cloud-powered advancements in storage management, hypervisor capabilities and layer-2 network extensions have enabled data to be replicated at great speeds, with greater fidelity and with greater reliability than ever before — enabling dramatically shorter RTOs with virtually no data loss.

Insight #6: aaS Delivery Model Unlocking Flexible, Practical Access to Next-Generation DR Solutions

While cloud-powered DR solutions enable new ways of achieving RTO and RPO improvements, the reality is that these improvements were always possible — they just came at a (often prohibitively) high cost. Today, budget pressures remain as high as ever, and the pursuit of optimal cost-effectiveness is still the de facto driver of many companies' DR programs. This is where cloud-based solutions enable a powerful synergy, bringing together the enhanced computing power of the cloud, while making powerful new capabilities eminently more accessible, scalable and cost-effective through the as-a-Service delivery model.

Making Warm- and Hot-Site DR Cost-Effective

The basic options for disaster recovery and high-availability programs remain the same:

Configuration	Pros	Cons
Cold Site	• No capital outlay if services contract option is used • Basic recovery option • Low telecom costs	• Capital outlay unless contract services option is used • Need self-discipline to test consistently • Time-consuming, high-risk restore process with higher rate of restore failure • Contract or systems must be synchronized with live system changes
Warm Site	• Provides periodic restore testing as part of solution • Provides backup to last data refresh date • Quicker recovery time than cold site • Low telecom costs	• Capital outlay unless contract option is used • Contract or systems must be synchronized with live system changes • Still relatively slow to restore • High service costs for periodic restores
Hot Site	• Near-real time recovery • Costs can be allocated to other functions (e.g. high availability) • Good development/ reporting environment • Reduced load on primary equipment	• Higher cost solution • Maintenance of 2 active environments • Telecommunication costs

The major change is that the new, cloud-based aaS model makes warm- and hot-site approaches dramatically more practical and cost-effective. Whereas in the past, the cost to set up a high-availability environment was at least equal (and often greater) to the cost of the production environment, leading cloud-based options have already brought that down to just 25-50% of the production costs. That is to say, the cost of achieving the highest standards of RPO and RTO has dropped by at least half.

The Hybrid Cloud Emerges as the Ideal Model: Flexibility + Control

Over the last few years, the business world has moved from a highly contested debate over the relative benefits of public vs. private clouds, to a consensus that the hybrid cloud model is the key to a cloud-powered digital transformation. Delivering the best of both worlds, the hybrid model enables scalable solutions that are built upon, and supported by, predictable foundations. Companies can take advantage of the flexibility of pay-as-you-go resources that are easy to deploy and scale up — while still maintaining the security and control they want and need.

The New Model: Disaster Recovery as a Service

Leading vendor partners in the disaster recovery and business continuity space are bringing together the most advanced, next-generation DR technologies — including cost-effective hot- and warm-site high-availability environments — in a smart, flexible as-a-Service delivery model. These DRaaS solutions include all of the infrastructure, process management and services necessary to restart an entire workload in a public cloud environment. All of the configurations are built in advance so that if necessary, failover can happen very quickly. The DRaaS model gives organizations the benefits of business continuity while removing the burden of configuring and maintaining hardware and software — and provides a trusted partner to assist in the recovery of your data in the event of a failure. In addition, organizations who are looking to move some of their operations to the cloud can take advantage of a scalable OPEX model while adding additional security through offsite services.

The True aaS Differentiator: Expertise + Consultative Service

An increasing number of vendors now offer cloud-based subscription models for DR software. But delivering technology through the cloud in a subscription model is only part of the equation. The true aaS model requires a robust service component. True DRaaS partners deliver comprehensive, consultative service that makes their expertise an extension of your team — enhancing your DR program while significantly alleviating the workload on internal IT and Security teams that are already overburdened. Leading DRaaS partners take on the end-to-end process of planning, building, testing and maintaining a DR program, from conducting an in-depth assessment of your business needs, creating DR strategies to achieve specific business goals, managing rapid implementation of a customized DR program, and executing thorough DR testing and ongoing optimization to keep your business ready for what's next. The result is a more effective, more reliable, higher-performance DR program — with lower monetary and internal staff costs.

The DRaaS Advantage

1. Flexibility: Customize your solution to meet your business continuity goals

RTO

How fast do you need to recover?

RPO

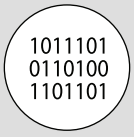

How much data loss can you withstand?

Budget

How much money do you have to spend?

2. Speed: Empower critical business agility

Rapid Deployment

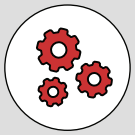

Implement a DR program and build an HA environment in weeks — not months or years.

Real-Time Scalability

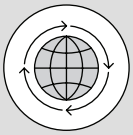

Expand your DR and HA services as your business needs evolve — and pay only for what you use.

3. Service: Leverage partners to get better results with less internal burden

Consultative Expertise

Align your DR program with your unique & evolving business goals.

Hands-On Support

Outsource end-to-end DR administration to a team of experts.

Conclusion: Stepping Up to Next-Generation DR

The 2020 coronavirus pandemic served as a wakeup call on disaster recovery and business continuity planning. In fact, most businesses proved remarkably resilient and agile in response to the disruptive challenges of the pandemic. Yet the acute experience forced many companies to recognize that their business continuity strategies needed to get better — not to prepare for the next pandemic, as much as simply to prepare for the much more common threats of hardware failure, user error, and local power and network interruptions. Businesses are growing more digital by the day, expanding their risk exposure to these everyday threats. And the stakes are also rising: Customers increasingly expect continuous, reliable service from every business, making the potential damage — in lost revenue, as well as lost customer trust, loyalty and brand reputation — greater than ever.

Of course, the flip side of these higher stakes is a major opportunity. The businesses that rise to the challenge, protecting business continuity and delivering reliable, continuous service, will meet and exceed customer expectations — and gain an invaluable competitive edge in their markets.

Capturing this opportunity undoubtedly requires more powerful technological capabilities — things like asynchronous disk replication, active backup, and hot-site, high-availability environments. And while these capabilities have traditionally been costly and difficult to implement, the emergence of the cloud-powered, subscription-based delivery model now gives organizations of all types and sizes practical, cost-effective access to next-generation DR capabilities. But, critically, for most organizations, subscription-based access to the technologies alone will not be enough — administering, testing and optimizing their DR programs will remain a complex, costly challenge and a barrier to their DR success.

Organizations looking to realize the promise of the DRaaS model must seek out solutions that deliver the synergy of technology and expert support within the aaS offering. This combination is the key to cracking the equation of improved RTO and RPO — at a lower cost. Moreover, it's the key to enabling the full value of a disaster recovery program in its simplest sense: Giving an organization the peace of mind of knowing its business operations, its valuable data, and its customer relationships are protected — no matter what happens.

Sources

1. The Real Costs Of Planned And Unplanned Downtime: Accelerate Recovery With New Technologies. Forrester Opportunity Snapshot: A customer study commissioned by IBM. August 2019. https://www.ibm.com/downloads/cas/L57KW7ND
2. 2019 IT Outage Impact Study by Logic Monitor. https://www.logicmonitor.com/it-outage-impact-survey/
3. Downtime costs small businesses up to $427 per minute. Mark Brunelli. October 08, 2015. https://www.carbonite.com/blog/article/2015/10/downtime-costs-small-businesses-up-to-$427-per-minute/
4. Cost of Data Center Outages. Data Center Performance Benchmark Series. Independently conducted by Ponemon Institute LL. January 2016. https://www.vertiv.com/globalassets/documents/reports/2016-cost-of-data-center-outages-11-11_51190_1.pdf
5. How Protected Is Your Data?. Kristi Rascon. March 2019. http://dellemcstudy.blogspot.com/2019/03/how-protected-is-your-data.html
6. Human Error is Top Cause of Downtime. Laura DiDio. October 01, 2018. https://ibmsystemsmag.com/Power-Systems/10/2018/Human-Error-Top-Cause-of-Downtime

RICOH
imagine. change.

www.ricoh-usa.com

Ricoh USA, Inc. 300 Eagleview Boulevard, Exton PA 19341 | 1-800-63-RICOH