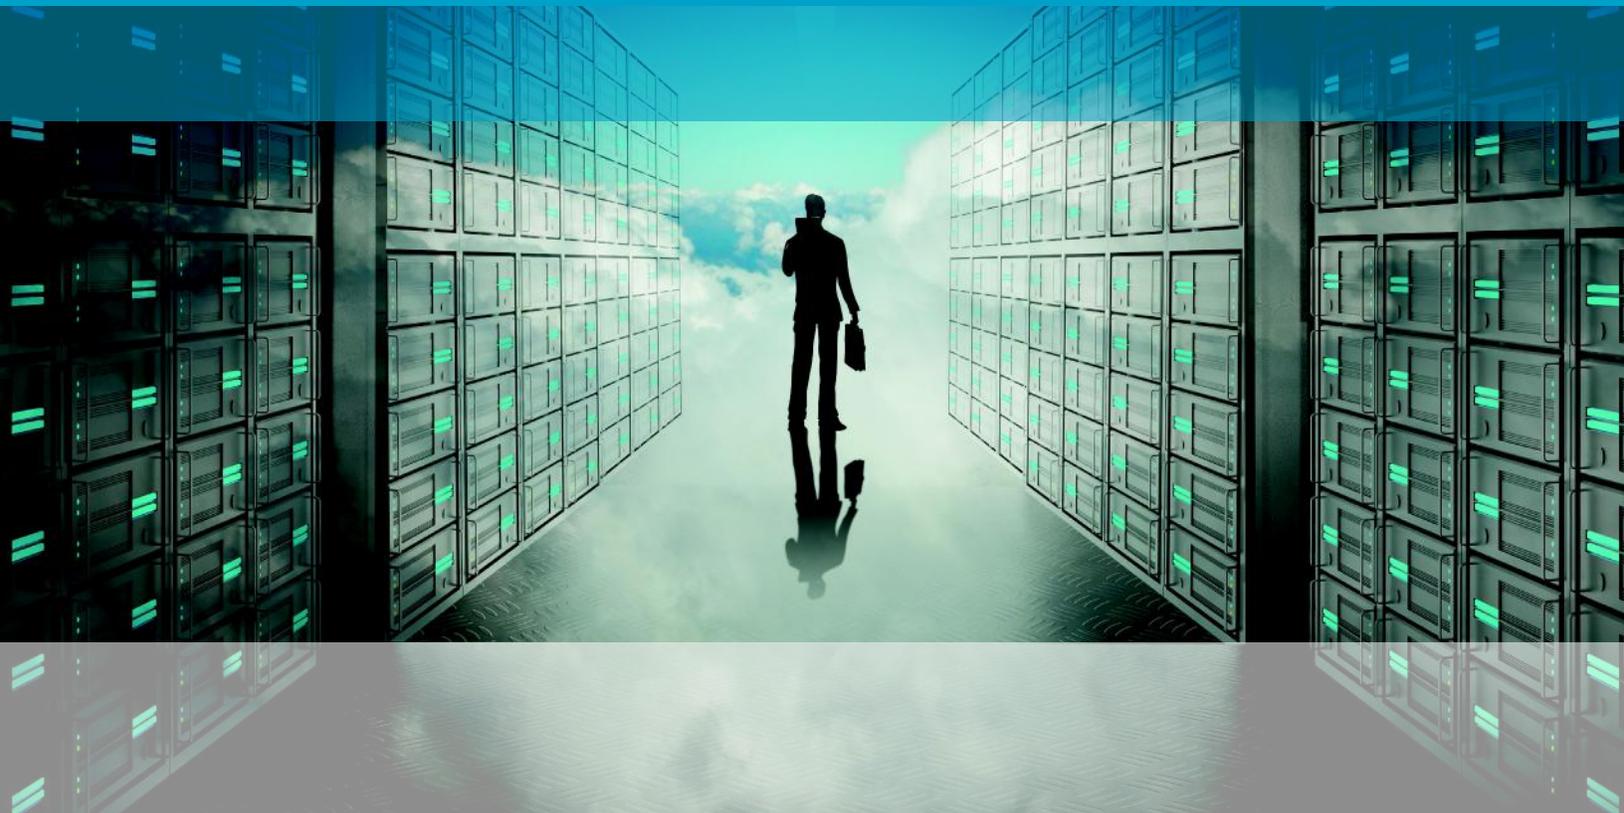


Recovery as a Service Comes of Age



Have security and availability concerns kept you from looking at backup and recovery in the cloud? The wait is over. Here's a short primer on the benefits and features of the new and improved recovery as a service.

By Derek Siler, Director of Solution Engineering, Sungard Availability Services; Mike McElaney, Director of Carrier Services, Presidio; and Alex Foster, Cloud Project Manager, Cisco Systems.

Increased comfort with cloud security. Avoidance of capital expenditures. Outsourced management. Pay-as-you-grow pricing. These and other features and benefits, combined with mounting concerns over maintaining business continuity in the event of disasters, are behind the growing adoption of recovery as a service (RaaS). This cloud-based approach offers service packages tailored to different levels of data and application criticality within organizations. Costs are based on speed of recovery and the restoration point for individual applications.

Recovery Becoming a Priority

The increased dependence on technology by organizations, and a widespread awareness of the high costs of downtime, are behind a booming market for disaster recovery services. Downtime is incredibly costly—from \$20,000 to \$200,000 per hour in labor and services costs, according to a 2012 study by Deloitte.¹ In some cases it can also be devastating, leading to a loss of business, a loss of customer confidence, and ultimately bankruptcy for some companies.

According to a 2013 report from ABI Research,² “As more critical infrastructure sectors like energy and healthcare go online, business continuity and disaster recovery become increasingly important.” The report attributes the adoption of cloud services for disaster recovery to “increasing global connectivity, faster communications networks, better technology, and the proliferation of high-end mobile devices” used for business inside and outside of the enterprise firewall.

Accelerating Adoption of Cloud Services

Recovery services in the cloud were available beginning several years ago. But many organizations were hesitant to trust the cloud for business continuity due to concerns about the security risks involved in replicating their mission-critical applications and data outside of the enterprise data center. These concerns encompassed cloud services in general, even as security, availability, and cloud business models continued to evolve and improve. But that has changed dramatically in recent years. A December 2012 report from IDC³ found that global revenue from hosted private cloud service providers was forecast to grow from \$3.15 billion in 2011 to \$24.3 billion in 2016. A 2013 IDC report⁴ uncovered even higher spending on public IT cloud services, which are forecast to grow from \$47.4 billion in 2013 to more than \$107 billion in 2017. And this comfort with cloud services is directly affecting the disaster recovery market, with a 2013 study by Markets and Markets⁵ finding that the global RaaS and cloud-based business continuity market is poised to grow from \$640.8 million in 2013 to \$5.77 billion by 2018, an incredible compound annual growth rate (CAGR) of 55.2 percent. A January 2014 report by Forrester⁶ echoes this forecast, finding that “the past few years have seen significant growth and adoption of RaaS across all segments, as infrastructure and operations professionals are looking for ways to improve their recovery objectives without increasing spend.”

¹ “Multivendor Network Architectures, TCO and Operational Risk,” Deloitte Consulting LLP, February 2012.

² “Cloud Security Services,” by Michela Menting, Senior Analyst, and Stuart Carlaw, Chief Research Office, ABI Research, April 9, 2013.

³ “Worldwide Hosted Private Cloud Services 2012–2016 Forecast: New Models for Delivering Infrastructure Services,” by Robert P. Mahowald, Frank Gens, Richard L. Villars, Satoshi Matsumoto, Chris Morris, Melanie Posey, Margaret Adam, Mark Schruft, Vladimir Kroa, Connor G. Sullivan, Ricardo Villate, Alejandro Florean, David Senf and David Tapper, IDC Research, December 2012.

⁴ “Worldwide and Regional Public IT Cloud Services 2013–2017 Forecast,” by Frank Gens, Margaret Adam, David Bradshaw, Christian A. Christiansen, Laura DuBois, Alejandro Florean, Phil Hochmuth, Vladimir Kroa, Robert P. Mahowald, Satoshi Matsumoto, Chris Morris, Tony Olvet, Kelly Quinn, Mary Johnston Turner, Richard L. Villars and Melanie Posey, IDC Research, August 2013.

⁵ “Disaster Recovery as a Service Market [RaaS, Cloud DR, Disaster Recovery as a Service, Business Continuity as a Service], Worldwide Forecasts and Analysis (2013 - 2018)” by Markets and Markets, March 2013.

⁶ The Forrester Wave™: Disaster-Recovery-as-a-Service Providers, Q1 2014,” by Rachel A. Dines, Forrester Research, January 17, 2014.

Today's Recovery Options

For midsize to large organizations, disaster and operational recovery are do-it-yourself activities, an outsourced service, or a combination of the two.

Do-it-yourself recovery is an expensive proposition requiring the high capital and operating costs of building and maintaining a separate data center at a remote location. Internal IT staff must be dedicated to this effort, reducing resources available for strategic business initiatives. As enterprise environments are continually changing, ensuring that all applications and data are 100 percent recoverable is an ongoing challenge requiring continual oversight. The speed of the recovery of the network (known as the recovery time objective [RTO]) and the precise point in an application's processing at which a recovery occurs (known as the recovery point objective [RPO]) are key measures in data and application recovery. Third-party professional organizations that specialize in disaster recovery are often much better at delivering speedier and more varied recovery options.

Another option is to contract for data center resources at a managed service provider's data center and manually back up data and applications to that infrastructure on a regular basis. This is more affordable, with shared physical infrastructure where data and applications from several companies may coexist on the same servers and storage repositories. But it does entail operational costs for an organization's IT personnel to administer. And for many companies, mission-critical applications and data are too sensitive for shared or multitenant environments.

A third option is a recovery in the cloud service, or RaaS. This service features a failover of a production application environment in an enterprise data center to a managed service provider's cloud environment. For the customer this managed service requires no capital expenditures, less complexity than do-it-yourself or hosted services, and more predictable monthly costs. Instead of having to move a primary data center to the cloud, the RaaS service replicates the on-premises IT environment for business continuity. There is no need to buy and maintain dedicated resources in case of a disaster. Instead, companies can purchase resources on a pay-per-use basis. Servers, storage, and network resources are available for use when they are needed, in an actual disaster or for testing. This pay-per-use model dramatically lowers costs, and RaaS plans offer guaranteed service-level agreements (SLAs) based on the RTO and RPO thresholds that have been agreed to.

Trimming Costs for RaaS by Prioritizing Recovery Times

All data and applications are not created equal. In other words, some organizations consider certain applications and data mission-critical (such as e-commerce or banking), whereas others, while important, are considered less critical to the organization's existence and reputation with customers in the short term (such as email and enterprise resource planning systems).

Costs for RaaS are highest for services in which the RPO is near zero (typically within 5 to 15 minutes) and the RTO is also short (ranging from 15 minutes to 4 hours). The most mission-critical applications also require an active/active clustered configuration with asynchronous server-based replication. It's managed end to end by the cloud provider to deliver recovery of applications in less than 4 hours and includes continuous data protection that enables restoration at any point in time within a 3-day timeframe prior to the outage.

Less mission-critical data and applications may include a near-zero RPO and an RTO of 2 hours or less. Or some applications may be serviced by disk-based backup with an RPO of 12 hours and an RTO of a full day. Applications that are much less mission critical can be served with tape backup failover and RPO and RTO within 72 hours. Thus, different applications and data can be assigned different storage media and RPO and RTO times based on their criticality to bring RaaS costs down.

What to Look for in a RaaS Solution

Many providers are now offering RaaS. How should organizations choose the best among them? Here are a few key things to look for:

- Make sure that the cloud provider can backup both physical and virtual infrastructure. Most cannot.
- Get a clear explanation of where data is stored and how it is handled throughout its lifecycle. Security should include the physical perimeter of a provider's data center, not just the network itself.
- Make sure the provider offers a guaranteed SLA that meets your organization's needs. There should be a financial penalty for failure to deliver on that SLA.
- The provider should have a detailed understanding of the compliance and regulatory issues that affect your organization.
- Change control should be up to date for the recovery production environments.
- The RaaS solution should be able to quickly and easily scale up or down in response to your needs.

For More Information

Explore the many benefits of RaaS for your company with a look at the world-class Recover2Cloud® solution from Sungard Availability Services (Sungard AS), brought to you by Presidio. Recover2Cloud is among the few recovery services for applications on an enterprise-class cloud infrastructure with guaranteed SLAs. What differentiates Recover2Cloud from other competitive solutions? Recovery for both physical and virtual environments. Strict adherence to SLAs. Data centers built on industry-leading converged infrastructure; including Cisco Unified Computing Systems (UCS), EMC Storage Arrays and VMware virtualization technologies. And Sungard AS has the capacity and expertise to recover a wide variety of IT environments, including legacy midrange and non-x86 systems. Sungard AS was recently cited as a Cisco partner in the creation of the world's largest global Intercloud – a network of clouds.

As IT experts, Presidio has helped more than 1000 companies design and implement enterprise-class data centers. Presidio can help you assess which applications are a best fit for Recover2Cloud and your company's readiness for a recovery as-a-service plan. For more information on Recover2Cloud, call 1-866-714-7209 or visit www.sungardas.com