

Cloud Recovery for Hybrid IT Environments

Recovery-as-a-Service (RaaS) leverages a number of infrastructure options and recovery platforms to provide a recovery solution that is flexible, scalable, and secure. Service providers can offer fully managed solutions with the expertise to support heterogeneous IT environments.

Application recovery on cloud-based infrastructures offers the benefits of shared tenancy — greater efficiency and lower cost — combined with the faster recovery speed and improved scalability that comes with virtualization

For all the benefits of cloud, the fundamentals of an effective recovery strategy remain the same, and essential.

Applications expertise is required to restore complex environments without delay and added downtime. Before selecting a recovery solution, organizations must analyze the value of their applications and the impact that application downtime would have on their business. Once those applications are analyzed, the work of prioritizing them into tiers, based on their criticality, remains important because it allows an IT organization to prioritize its investments in recovery technology. Finally, in every organization, the need to maintain and test recovery procedures is constant.

How do you take your cloud?

Cloud is notoriously difficult to define because cloud capabilities are not black-and-white. And most organizations are in the process of evolving their environments to a cloud-based architecture. Consequently, there are many ways to implement cloud-based recovery.

Determining the appropriate cloud implementation for your organization is a business decision based on two criteria: cost and security. The key is to look for ways to manage the trade-offs between the two. Greater security usually comes at greater cost, while a higher degree of automation is often tied to a more public cloud infrastructure.

Having the flexibility to implement cloud in a way that matches your organization's specific business needs requires the ability to choose from and blend multiple cloud and virtual technologies.

“Organizations must analyze the value of their applications and the impact that application downtime would have on their business.”



When a service provider assumes responsibility for recovery, RTO and RPO requirements can be backed by service level agreements (SLAs).

Never compromise on reliability

Regardless of your specific cloud implementation — the blend of technologies supporting your application recovery — the one thing that cannot be neglected is the reliability of the recovery solution itself.

Any successful recovery, cloud-based or otherwise, has a number of requirements:

- **Coherent data**
- **Sufficient capacity**
- **Available staff**
- **Time — for configuration, backups, and testing**
- **Recovery expertise**
- **Procedures that are both sufficiently tested for effectiveness and kept current with changes to the production environment**

To ensure these requirements are met, many organizations opt to have their recovery fully managed by an external service provider.

When recovery is cloud-based, additional questions that an organization must consider — and be sure to ask their provider — include:

- **Who owns the compute assets?**
- **Who owns the responsibility for recovery?**
- **Is there a Recovery Assurance associated with the service?**

Because a service provider has the resources and expertise in place to manage the full recovery, it can (and should) assume ownership of managing the compute assets and maintaining them to ensure the requirements for application recovery are met. In the best case scenario, the service provider will be responsible for all of the steps in the recovery lifecycle, including procedure planning, implementation, testing, and the recovery itself.

When a service provider assumes responsibility for recovery, RTO and RPO requirements can be backed by service level agreements (SLAs).



“Lower cost and scalability are the fundamental benefits of moving recovery to the cloud.

But for critical applications, the service provider should be selected based on security and reliability.”

The building blocks of a reliable cloud-based recovery

Developing a reliable cloud-based recovery strategy calls for a tiered approach to application availability. The first step is to perform a business impact analysis (BIA). This exercise defines which business processes are most critical — and most impacted by application downtime. Based on that analysis, applications are prioritized and organized into tiers based on their sensitivity to downtime. Each tier can then be defined by the recovery time and recovery point requirements of the applications within it.

Modernized data movement technologies

Every data movement technology has an associated RTO/RPO, and these application tiers present an opportunity to modernize data movement to meet the RTO and RPO requirements of specific applications. Selecting the right data movement technology for each tier is essential for an organized, efficient and successful recovery. For example, traditional tape-based backup delivers recovery anywhere from two days to a whole week. For critical applications like e-mail and online order systems, that level of downtime could shutter many businesses for good. Online technologies, on the other hand, can help get applications back online in a matter of hours.

Three categories of technology modernize data movement to best support the full range of application recovery requirements: **server replication**, **storage replication**, and **vaulting**.

Server replication is ideal for critical applications in environments with a smaller numbers of servers. **Storage replication** is also a right solution for critical applications, and is best deployed in environments with larger numbers of servers. **Vaulting**, also known as online backup and recovery, eliminates the cost and delay of tape-based backup.

Selecting a service provider

Each of these solutions can be deployed in the cloud. But the question remains: When selecting a cloud-based service provider, what criteria should be applied?

Lower cost and scalability are the fundamental benefits of moving recovery to the cloud. But for critical applications, the service provider should be selected based on security and reliability.

Fully-managed recovery services backed by contractually-guaranteed RTO and RPO service levels are the key to assuring application availability. Additionally, the ability to select from a full range of technologies for moving data to the cloud will ensure applications are restored in priority order. Finally, a full range of infrastructure options — including managed, self-managed and managed private cloud offerings — eliminates over spending.



“Every IT application’s environment has specific recovery requirements.”

Managed cloud-based recovery

Every IT application’s environment has specific recovery requirements. The following model illustrates how a managed solution applies a range of technologies to help ensure recovery requirements are met in a way that ensures security, reliability, and cost-effectiveness.

Selecting the technology

With cloud-based recovery, once the BIA is completed and tiers are defined, applications and data can then be protected on a shared, secure cloud using the most appropriate data movement technology.

For the most critical applications, server replication or storage replication solutions are selected for their ability to deliver recovery in less than four hours at a near-zero recovery point.

Server replication software encapsulates the OS, applications configurations and applications data, and replicates them as soon as changes are written to disk. Similarly, in bigger environments utilizing storage replication, software deployed on the storage array cost-effectively replicates large numbers of virtual machines (VMs).

Those applications that are less sensitive to downtime but no less critical to the business are encrypted and stored in a secure offsite location using vaulting technology for online backup and recovery. Vaulted data can be restored to the last successful backup in less than 24 hours.

Handing over recovery responsibility

With a managed solution, the service provider is now responsible for the management and maintenance of the full recovery lifecycle.

During normal operation, the service provider executes 24/7 monitoring and handles monthly capacity tuning.

For testing, the service provider delivers recovery runbooks customized to your environment and supplies end-user verification as part of the exercise. When the service provider’s staff does the heavy lifting required for recovery testing, it eases the burden on your IT staff and can eliminate the travel and expenses normally associated with a test.

In the event of an outage, disaster, or other interruption, the service provider will recover applications and data into VMs — and promote those VMs as active systems. Automating as many steps in the process as possible speeds the recovery and avoids human errors and human dependencies. The replication technologies provide ready-to-operate VMs and data standing by at the recovery site while the restoration and rebuilding of VMs from vaulted data takes place.

The service provider coordinates the recovery in concert with additional applications on physical systems and legacy mainframe equipment, and reconnects business users to the recovered applications and data. When the incident has passed, the production environment is then returned to normal operations.



Cloud Recovery for Hybrid IT Environments

Conclusion

Managing recovery effectively is complex. A fully managed recovery solution removes the burden of that complexity from an IT organization. And when the recovery solution leverages modernized data movement technologies and a range of secure cloud and virtual infrastructures, the advantages are significant: higher availability, better performance, and lower cost.

What to look for in a cloud-based recovery solution:

Managed cloud services

- Service provider is responsible for ensuring
- recovery of applications
- Recovery backed by contractually-guaranteed service-level agreements (SLA)
- SLA covers the RTO/RPO requirements of your applications

A full range of data movement capabilities

- Ensures RTO/RPO match the business value of applications
- Prevents overspending and under-protection

A full range of infrastructure options

- Self-managed cloud
- Managed cloud
- Managed private cloud

Enterprise-class infrastructure

- Actively managed and maintained
- Ensures security and reliability

For more information please visit our website at:

www.sungardas.com/DR

About Sungard Availability Services

Sungard Availability Services provides managed IT services, information availability consulting services, business continuity management software, and disaster recovery services.

To learn more, visit www.sungardas.com or call 1-888-270-3657

Trademark information

Sungard Availability Services is a trademark or registered trademark of SunGard Data Systems or its affiliate, used under license. The Sungard Availability Services logo by itself is a trademark or registered trademark of Sungard Availability Services Capital, Inc. or its affiliate. All other trademarks used herein are the property of their respective owners.

Connect with Us

