

VERITAS™



Veritas and Accenture: High-Availability and Disaster Recovery Scenarios

Using best-of-breed technologies to solve organizations' multi-cloud and hybrid challenges.

Contents

Introduction	3
Scenario 1: Ensure Multi-Cloud DR with InfoScale	3
Context and Client Challenges	3
Our Approach	4
Conclusion	7
Scenario 2: Build Hybrid-Cloud DR with Resiliency Platform	7
Context and Client Challenges	7
Our Approach	8
Conclusion11

INTRODUCTION

A hybrid cloud approach enables an organization to combine best-of-breed technologies while increasing cost-effectiveness and ensuring independence from a single hosting provider. There are many reasons why organizations are moving to the cloud, such as to:

- Become more agile and flexible.
- Innovate faster.
- Shift everything to the “as-a service” model.
- Simplify IT operations.
- Improve application high availability (HA).
- Enable low-cost disaster recovery (DR).

Thinking about an HA/DR strategy, implementing it and confirming it works is crucial to ensuring an organization’s business continuity is impacted by minimal application downtime. In this context, it’s essential to understand the ecosystem of the organization’s applications and their dependencies.

Typically, there are two business scenarios we encounter:

- **Multi-Cloud**—Organizations that are migrating their entire application services portfolio to the cloud or have adopted an all-in or “full-cloud” strategy with their cloud providers and need better application HA performance and resiliency.
- **Hybrid**—Organizations that prefer to adopt a hybrid model where application services are hosted on-premises and in the cloud or between clouds.

Veritas and the Accenture Cloud Innovation Center (ACIC) have worked synergistically to develop established blueprints to address these scenarios to support organizations during their journey to the cloud:

- **Multi-Cloud**—Using Veritas InfoScale™, we help organizations design an application-aware solution for HA and DR in a full-cloud context.
- **Hybrid**—Using Veritas Resiliency Platform, we help organizations during their migration to the cloud, managing hybrid DR architectures able to run nondisruptive tests and Instant Recovery Restores in a single pane of glass.

SCENARIO 1: ENSURING MULTI-CLOUD DR WITH INFOSCALE

Context and Client Challenges

Hybrid clouds are becoming standard for many enterprises. The factors driving this decision include increased flexibility as well as savings in capital expenses (CapEx) and operating expenses (OpEx). Although there is considerable interest in hybrid-cloud deployment models, ensuring seamless data movement, enterprise performance, scalability and resilience can be challenging. An unplanned disruption to the business/ data center can have highly destructive and far-reaching impacts on productivity, customer services, the supply chain and revenues. To avoid this outcome, enterprises must implement robust solutions to ensure their services are always up and running.

When it comes to the cloud, and specifically to Amazon Web Services (AWS), intelligently using availability zones and geographical areas can make the difference in business continuity. Such a strategy protects organizations against the failure of a single availability zone and/or geographical area, keeping the service always on. (See Figure 1.)

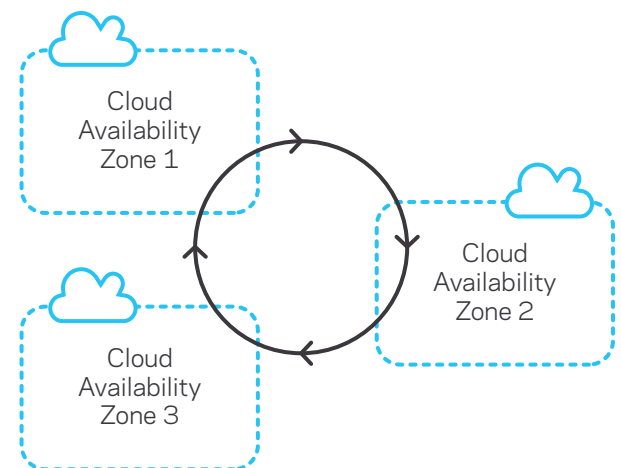


Figure 1. Enable application high availability (HA) by using multiple AWS availability zones.

In this context, InfoScale enables organizations to protect their most important applications from planned and unplanned downtime while providing HA and DR solutions for critical business services, including individual databases, custom applications and complex, multi-tiered applications across virtual and physical environments over any distance.

InfoScale monitors applications and infrastructure to detect possible risks to availability and automatically recovers applications when needed. InfoScale also detects site outages and initiates application recovery at a DR site. With built-in testing to proactively detect and report potential problems before they can impact IT services, InfoScale helps minimize unwanted downtime of critical business services.

Our Approach

To address the need for a service to be always up and running, Accenture and Veritas chose Veritas InfoScale technology. As software-defined storage solution, InfoScale offers enterprise-grade capabilities for organizations looking to deploy and run top-tier applications across multiple clouds. InfoScale delivers a complete suite of capabilities for organizations looking to deploy a robust, multi-cloud infrastructure that includes public clouds such as AWS, Microsoft Azure and Google Cloud. InfoScale does so by offering:

- **Intelligent data movement to the cloud**—Organizations need to ensure they're moving data to the cloud in an optimal way. InfoScale includes SmartMove technology that ensures efficient, intelligent data migration to help organizations save on bandwidth and storage costs while migrating application data to the cloud. SmartMove analyses storage usage by looking up the file system (VxFS) metadata and moving only relevant or active data to the cloud. For example, if an organization has provisioned 10 TB of storage and has 3 TB of active (non-deleted) data, then a standard data mover typically moves 10 TB of data to the cloud by considering all non-zero-bit strings as valid data. In contrast, InfoScale's SmartMove technology looks up the file system metadata and then moves only the non-deleted or active data to the cloud. SmartMove can also optimize data movement to the cloud in terms of CPU, bandwidth and schedule.
- **Applications failure management in seconds**—Many, if not most applications require HA in the event of a node (compute and/or data), availability zone or regional outage. InfoScale provides automated remediation actions via its Intelligent Monitoring Framework (IMF), which monitors applications at a service level and triggers a failover in seconds, not minutes, when it detects an application failure. Organizations can even perform a "fire drill" to identify potential problems and ensure applications will failover predictably. InfoScale offers application-aware agents for leading enterprise applications, and organizations can write custom agents for in-house applications. Volume and file-level replication ensure application data is replicated to other AWS availability zones or regions to protect against a large-scale infrastructure outage.
- **Software-defined storage capabilities**—Today's data centers have multiple layers of physical and virtual environments, each with its own administrative tools, with applications or entire business services composed of multiple components that interact with each other despite running on different operating systems and virtualization technologies. InfoScale provides automated, orchestrated recovery of applications composed of multiple components running on different virtual tiers (see Figure 2).

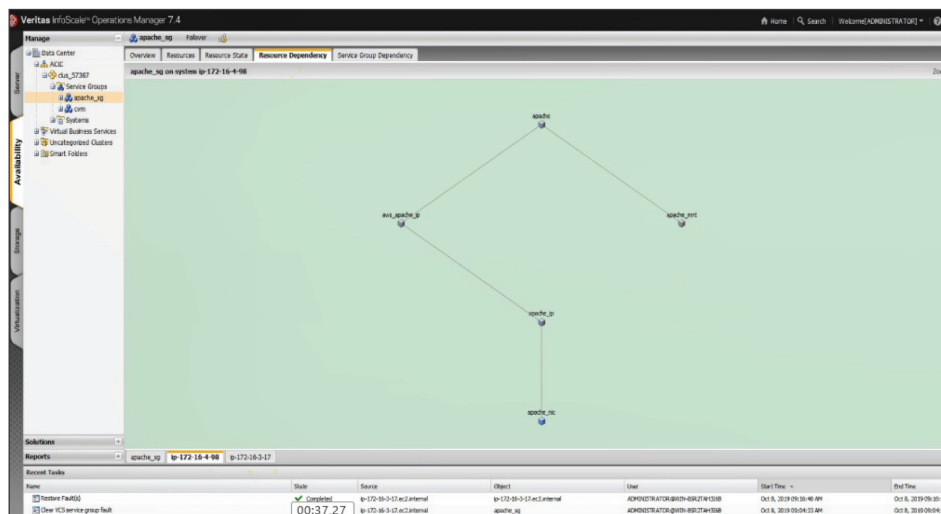


Figure 2. Monitoring Apache application components via InfoScale.

With InfoScale, IT administrators can not only recover failed applications but also automatically orchestrate the connection to other computing resources needed to keep the business service available via a faster recovery with minimal downtime and no manual intervention. (See Figure 3.)



Figure 3. Admins can use InfoScale's central dashboard, Operations Manager, to manage data and application failures.

- **Compatibility with hundreds of applications and databases**—InfoScale provides off-the-shelf support for a wide range of applications, reducing deployment times for SAP®, BEA®, Siebel®, Oracle applications, Microsoft Exchange and PeopleSoft as well as enterprise-class databases such as Oracle, DB2, Microsoft SQL Server and Sybase.
- **Storage reliability and performance**—Mission-critical applications demand the same performance and resiliency characteristics as they do on-premises. The SmartIO and Flexible Storage Sharing (FSS) technologies in InfoScale provide performance gains and shared storage capabilities that deliver predictable service-level agreements (SLAs).
 - **SmartIO**—In AWS, SmartIO uses instance store-SSD storage closest to compute as a data caching device to improve performance. The result: InfoScale's SmartIO technology accelerates the workload's throughput by 4X (for a 20 percent cache hit).
 - **Flexible Storage Sharing**—FSS in AWS uses elastic block storage (EBS) volumes and creates a shared-nothing cluster to meet protocol flexibility and performance needs and delivers horizontal scalability for enterprise applications. The result: InfoScale's FSS technology allows organizations to scale IOPS linearly. Organizations can even run multiple applications in a single FSS cluster and set up SLAs to ensure isolation among applications. This type of deployment unlocks valuable use cases such as running real-time analytics on incoming transactional data or fraud detection on credit card transactions.
- **Broad hardware and platform support**—InfoScale is the only solution that supports all leading operating systems, including Windows, Linux and virtual platforms including VMware® ESX®, Red Hat® Enterprise Virtualization (RHEV), Oracle VM and Hyper-V. As proved in the ACIC Lab, InfoScale also ensures application availability in full-cloud environments hosted in AWS, Azure and Google.
- **Simplified migration and integration**—InfoScale monitors applications and, in case of failure, handles their recovery to a healthy virtual machine (VM) running on a different platform. InfoScale also enables organizations to easily migrate workloads, automatically supporting physical, virtual and all major public clouds, thus avoiding vendor lock-in. (See Figure 4.)

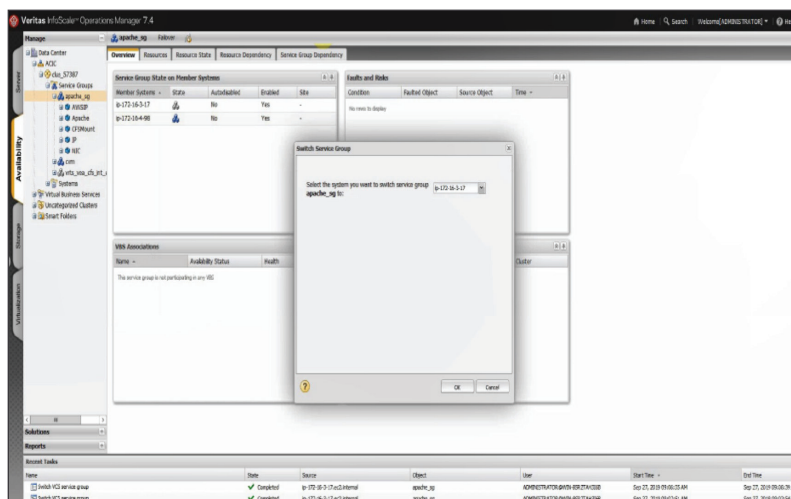


Figure 4. Monitoring application movement via InfoScale's central dashboard, Operations Manager.

In this specific ACIC scenario, we focused on an AWS full-cloud environment with application HA guaranteed by InfoScale and applied to multiple availability zones (see Figure 5).

Deployment of the InfoScale environment on AWS is based on the following best practices:

- Using AWS Cloud Formation to create cloud-based InfoScale environments to accelerate deployment (see Figure 6).
- Developing applications on AWS Elastic Compute Cloud (EC2) instances, not relational database service (RDS) instances.
- Using AWS Elastic Block Storage (EBS) to enable data mirroring between availability zones and data replication between geographical areas.
- Using AWS Elastic Load Balancer (ELB) to manage traffic routing between availability zones.
- Applying the required routing tables to allow application services to properly communicate between availability zones.
- Using a Coordination Point Server on a third availability zone (optional and not included in this scenario) to arbitrate in case of network failures between availability zones.

Following these guidelines and using InfoScale, we've enabled a high-performance, Apache-based application to be highly available, thus guarantying service continuity between different AWS availability zones in two main situations:

- Application downtime.
- AWS availability zone downtime.

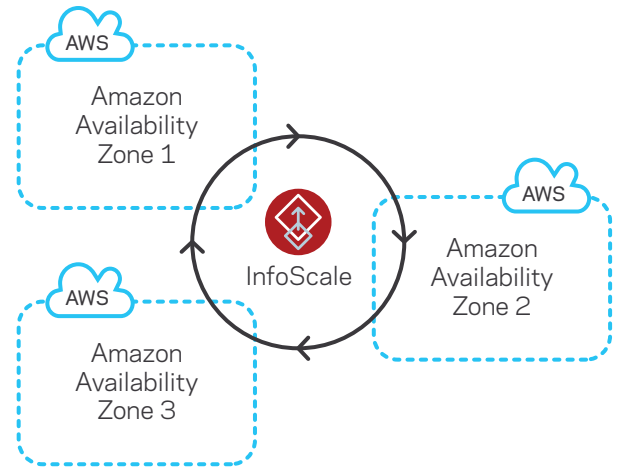


Figure 5. InfoScale enables application HA in a multi-availability zone scenario.

Filter by:	Status	Search events
2019-08-09		
16:17:47 UTC+0200	CREATE_COMPLETE	AWS::CloudFormation::Stack UC3INFOSCALEVERITAS
16:17:45 UTC+0200	CREATE_COMPLETE	AWS::AutoScaling::LifecycleHook TerminateLifecycleHook
16:17:44 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::LifecycleHook TerminateLifecycleHook
16:17:44 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::LifecycleHook TerminateLifecycleHook
16:17:44 UTC+0200	CREATE_COMPLETE	AWS::AutoScaling::LifecycleHook LaunchLifecycleHook
16:17:43 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::LifecycleHook LaunchLifecycleHook
16:17:43 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::LifecycleHook LaunchLifecycleHook
16:17:39 UTC+0200	CREATE_COMPLETE	AWS::AutoScaling::AutoScalingGroup AutoScalingGroup
16:16:34 UTC+0200	CREATE_COMPLETE	AWS::EC2::Route AZ2NATRoute
16:16:31 UTC+0200	CREATE_COMPLETE	AWS::EC2::EIPAssociation EIPAssociation
16:16:21 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::AutoScalingGroup AutoScalingGroup
16:16:20 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::AutoScalingGroup AutoScalingGroup
16:16:19 UTC+0200	CREATE_COMPLETE	AWS::EC2::Route AZ1NATRoute
16:16:19 UTC+0200	CREATE_IN_PROGRESS	AWS::EC2::Route AZ2NATRoute
16:16:18 UTC+0200	CREATE_IN_PROGRESS	AWS::EC2::Route AZ2NATRoute
16:16:18 UTC+0200	CREATE_COMPLETE	AWS::AutoScaling::LaunchConfiguration LaunchConfiguration
16:16:16 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::LaunchConfiguration LaunchConfiguration
16:16:16 UTC+0200	CREATE_IN_PROGRESS	AWS::EC2::EIPAssociation EIPAssociation
16:16:16 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::LaunchConfiguration LaunchConfiguration
16:16:15 UTC+0200	CREATE_IN_PROGRESS	AWS::AutoScaling::LaunchConfiguration LaunchConfiguration
16:16:15 UTC+0200	CREATE_IN_PROGRESS	AWS::EC2::EIPAssociation EIPAssociation
16:16:14 UTC+0200	CREATE_COMPLETE	AWS::EC2::NATGateway AZ2NATGateway
16:16:11 UTC+0200	CREATE_COMPLETE	AWS::EC2::Instance Bastion
16:16:03 UTC+0200	CREATE_IN_PROGRESS	AWS::EC2::Route AZ1NATRoute

Figure 6. AWS Cloud Formation deployment of InfoScale environments.

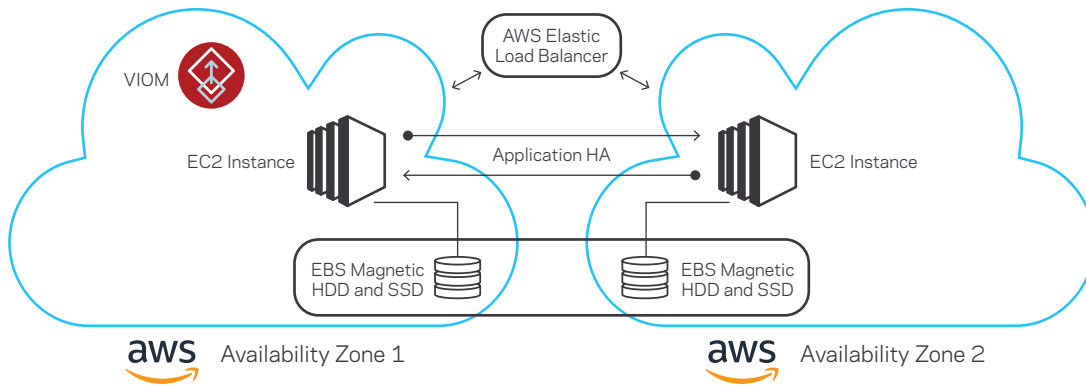


Figure 7. A high-level design of how InfoScale enables application HA/migration between AWS availability zones.

This scenario includes use of additional AWS features such as ELB to enable service continuity transparently (see Figure 7).

Conclusion

Organizations need to use new, innovative solutions and technologies to enable high-performance applications to become highly available and guarantee minimum business downtime. Accenture and Veritas are excited to fill that need by using InfoScale to provide the following advantages:

- Intelligent data movement to the cloud.
- HA between AWS availability zones and DR between different geographical areas.
- Storage reliability and performance.
- Application services that are always up and running by managing failures transparently to the application service.
- Software-defined storage capabilities.
- A simple approach for all the business-critical applications that supports a broad range of architectures: physical to virtual, virtual to virtual, physical/virtual to cloud, cloud to cloud and cloud to physical/virtual.
- Integration with Veritas Resiliency Platform to provide a single pane of glass for orchestrating any kind of architecture.

SCENARIO 2: BUILDING HYBRID-CLOUD DR WITH RESILIENCY PLATFORM

Context and Client Challenges

Business continuity orchestration is an essential part of any organization's planning and operation, and many are increasingly faced with the need to reduce infrastructure and management costs while deploying high-performance resilient architectures. This mandate translates into increased time and complexity required for proper DR management.

The cloud has been the next big thing in IT since virtualization, and organizations around the world are adopting it. Although the cloud offers big benefits with the potential for high business returns in the short term, organizations need to stay agile and view the cloud as just one strategic element of an independent business strategy that protects business applications around-the-clock and ensures a strong return on investment (ROI), regardless of where applications are hosted. The key is for organizations to be able to move business services on demand seamlessly to and from any single cloud or even across multiple, disparate clouds as needed. Such organizations need help managing and orchestrating the business continuity of their services as well as managing migrations from on-premises to the cloud and/or between clouds without losing the benefits of a hybrid-cloud strategy. (See Figure 8.)

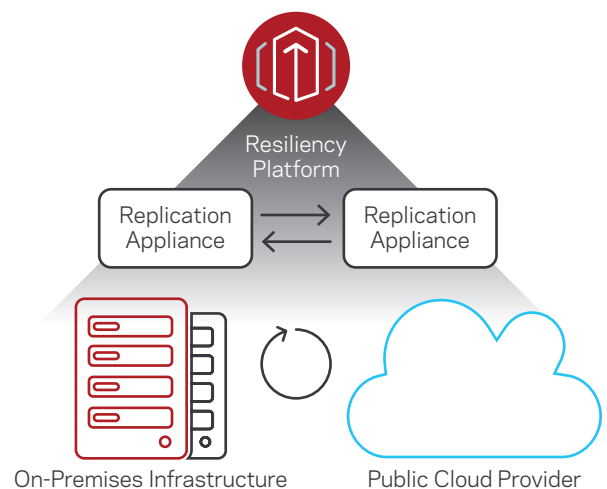


Figure 8. Enable easy management of a hybrid-cloud environment via a single pane of glass with Resiliency Platform.

Whether you're adopting the cloud as a permanent location for business applications or as a temporary home for applications in case of a DR scenario, there are several factors to consider before making the move.

- **Cloud migration**—Migrating business workloads to the cloud can be tricky and complex due to multiple moving parts. Organizations need to secure business data during the transfer from on-premises to the cloud and consider the dependencies within complex, multitier applications. Another potential roadblock is not having the ability to test migrations or to test how workloads perform in the cloud before moving services off-premises. Without advance testing, organizations can incur thousands or even millions of dollars in lost revenue if a migration doesn't go smoothly or if applications don't work as expected in the cloud. The biggest risk, however, arises when an organization hasn't carefully thought out its long-term cloud strategy. Today, the move to the cloud may be economical, but in a few years, an organization may want to move applications back on-premises or to another cloud if business priorities change. Without a failsafe for doing so, they may end up harming their long-term ROI.
- **Cloud disaster recovery**—DR is essential, and whether organizations are adopting a cloud-based DR approach for all business operations or just some of them, the potential CapEx savings can be tremendous. It's essential to ensure that OpEx doesn't suddenly skyrocket, however, organizations commonly employ application or workload resiliency measures to maintain business and application continuity. They could be managing resiliency operations across a hybrid-cloud architecture where resiliency personnel are managing legacy applications in on-premises data centers as well as managing resiliency across new cloud architectures, which requires additional training. Using multiple tools across different sub-environments can also lead to fragmentation, causing visibility and subsequent management issues and leading to an increased risk of downtime. Another key aspect many organizations fail to consider is whether they have a reliable way to test their DR strategy. Unfortunately, some still rely on an inefficient, all-hands-on-deck manual approach and weekend testing, which can be costly and time-consuming.

Our Approach

Despite today's often-complex application infrastructures, it's possible to ensure predictable resiliency for your business applications across hybrid and multi-cloud architectures. To do so, your business likely requires:

- A unified, automated approach that considers resiliency for applications across all locations, including the cloud.
- The ability to proactively test and re-test your migration strategy and DR preparedness easily whenever required without disrupting normal business operations.
- The flexibility to move your business applications to any target location, including the ability to migrate or failback applications from the cloud to on-premises at a moment's notice.

Organizations can finally address all these requirements with Veritas Resiliency Platform. (See Figure 9.)

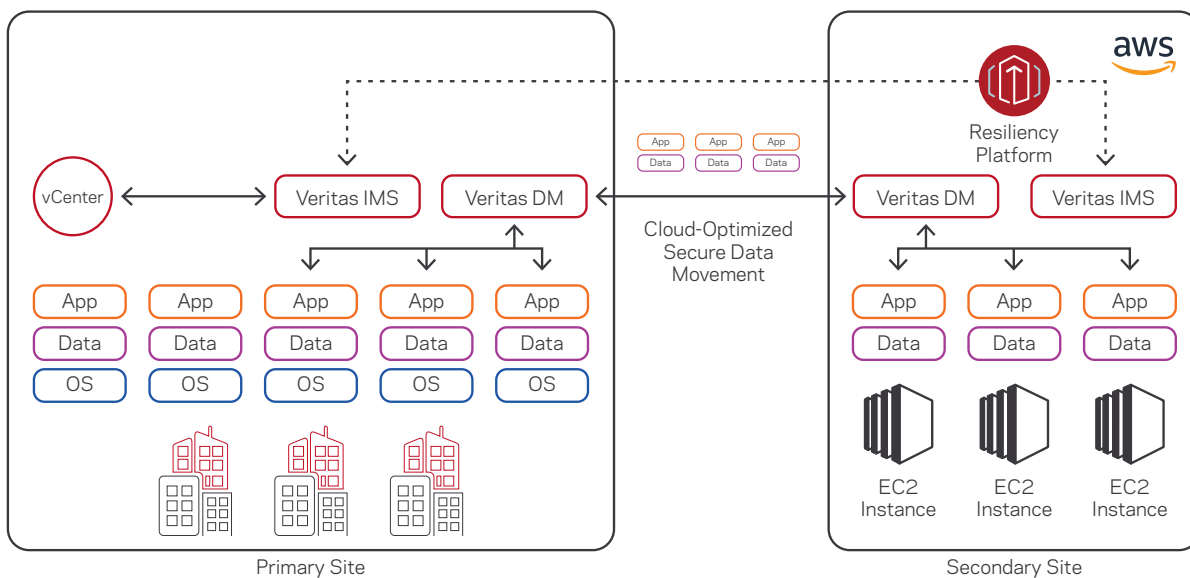


Figure 9. A high-level design of how to manage application DR/migration in AWS with Resiliency Platform.

Furthermore, integrating an organization's hybrid environment with Veritas NetBackup™ ensures a smoother recovery for all those VMs that don't need a short recovery point objective (RPO). The result is a ready-to-use environment with no waiting for restore time.

The key components of Resiliency Platform deployed in this scenario are:

- **Resiliency Manager**—Deployed as a virtual appliance hosted in the AWS cloud, Resiliency Manager is responsible for data center discovery, providing information on the assets to protect (storage arrays, VMs, applications) within a resiliency domain (the organization's infrastructure).
- **Infrastructure Management Server**—Deployed as virtual appliances hosted in the AWS cloud, Infrastructure Management Servers are exploited by the Resiliency Manager to discover and monitor assets (storage arrays, VMs, applications) within a data center to enable proper management and protection of the organization's infrastructure.
- **Data Mover**—Deployed as a virtual appliance hosted in the AWS cloud, the Data Mover is responsible for replicating VM data and configuration from the local (source) data center to the remote (target) data center. This capability addresses migration to the cloud as well as hybrid or full-cloud DR, including nondisruptive DR testing.

Using Resiliency Platform enables organizations to:

- **Orchestrate complex recoveries with predictability**—Organizations typically have large groups of applications and/or VMs, some or all with complex interdependencies that must be considered during stop and start procedures. When large-scale failures or site outages occur, organizations don't have time to waste on manual procedures. Their DR procedures must be executed quickly and flawlessly regardless of any underlying complexities. Ensuring success requires taking minimal steps along with providing a dashboard view to ensure visibility into when services come back online and identify any problems. Resiliency Platform fully automates all recovery procedures across a multi-vendor, multi-platform environment so admins can easily perform site-to-site large-scale recoveries with just a single click. A central dashboard provides visibility into how applications are meeting defined service-level objectives (SLOs) along with end-to-end visibility into recovery procedures to help monitor progress and resolve problems quickly, if required. (See Figure 10.)

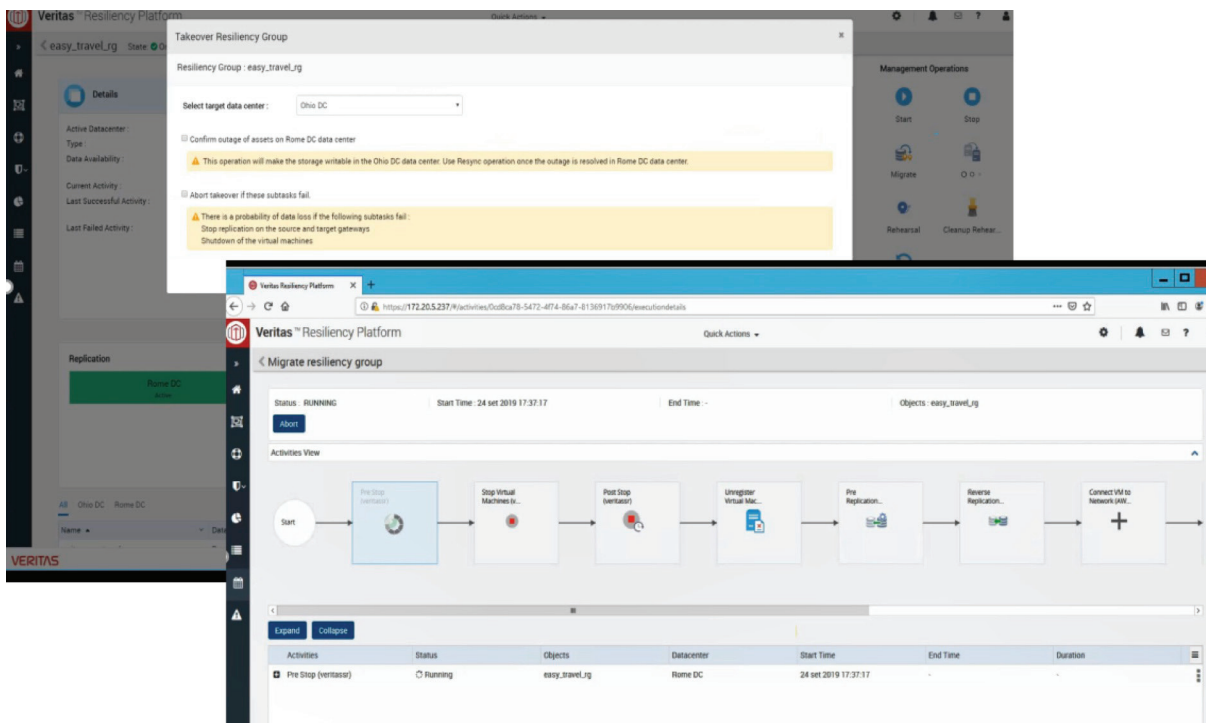


Figure 10. Easily manage a DR event via Resiliency Platform's central dashboard.

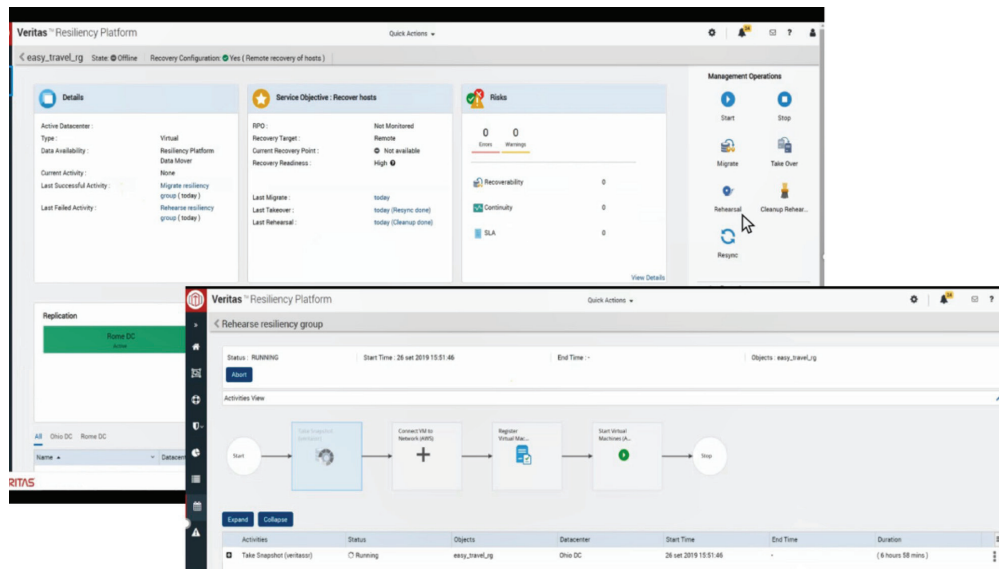


Figure 11. Execute a DR preparedness test with no impact to already-running applications with Resiliency Platform.

- Ensure DR readiness with nondisruptive testing**—Organizations shouldn't wait for a disaster to find out if their recovery plan will succeed. DR testing is vital, but is too often labor-intensive and carried out only during weekend hours. Businesses need to provide proof or assure internal and external stakeholders they will recover if a disaster is declared. They need the ability to test at any time without disrupting normal business operations and without requiring all-hands-on-deck during weekend hours. With Resiliency Platform and NetBackup, organizations can now nondisruptively rehearse recovery plans at any time to ensure proactive DR readiness and predictable recovery. Veritas provides a unified approach to testing that considers the need for business applications to recover data from disparate sources such as NetBackup or via hardware-based or host-based replication. The Veritas unified resiliency solution provides up-to-a-millisecond view of the time required for a business to recover services, performs automated clean-up to quickly release resources without manual involvement and enables organizations to prove compliance with internal and external business continuity mandates. (See Figure 11.)
- Achieve any SLO**—An organization's business applications likely span a range of SLOs. Backup and replication technologies are designed for different service levels and are usually managed separately with different processes and teams. Centrally managing resiliency for all SLOs requires centralized visibility and control of data availability across both backup and replicated data transports. Resiliency Platform's direct integration with NetBackup delivers resiliency for any SLO—from the top tier to the bottom—with simple management from a single console. This approach reduces manual efforts and OpEx, letting backup and virtualization admins concentrate on more important tasks. And most important, it helps avoid the chaos many organizations face without this level of unified resiliency.
- Seamlessly migrate complex, tiered applications between on-premises and the cloud**—Business applications shouldn't be tied to a single location. Instead, organizations should be able to move applications and data easily anywhere, depending on business and budgetary needs, whether it's to the cloud or back on-premises (see Figure 12).

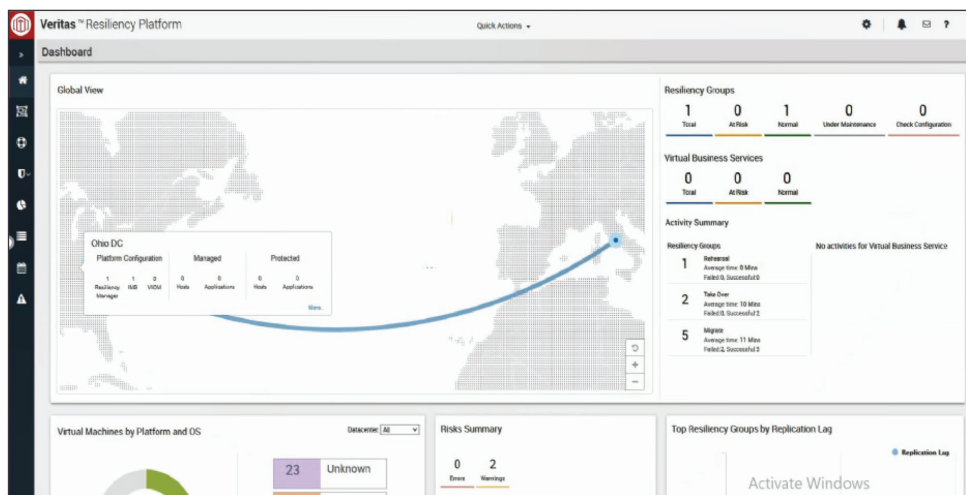


Figure 12. Enable easy workload migration in a hybrid-cloud scenario with Resiliency Platform.

- Resiliency Platform enables seamless movement of even the most complex workloads between on-premises and the cloud with a single click as well as quick migration of workloads including I/O-intensive applications to the cloud. Plus, businesses can stay agile and flexible with the option to failback services to on-premises, as needed.
- **Extend on-premises resiliency to the cloud with a unified hybrid-cloud strategy**—Don't let an on-premises workload and application resiliency strategy keep your organization from moving to the cloud. Instead, confidently expand from on-premises to multi-cloud architectures with a resiliency plan that scales easily and cost-effectively as the organization innovates while also ensuring application environments don't become fragmented. With Resiliency Platform, organizations can easily adopt the cloud with a resiliency strategy that cost-effectively scales from on-premises to the cloud and eliminates the need for multiple point tools across different parts of the environment. Choosing a single resiliency solution that spans on-premises, legacy environments and the cloud saves on OpEx, avoids environment fragmentation, increases visibility and minimizes the risk of downtime. This approach delivers real-time and historical SLO monitoring and reporting across an entire hybrid environment along with single-pane-of-glass visibility, so admins stay informed about the organization's IT health. Plus, IT only needs to be trained on one, easy-to-use, unified resiliency platform interface.

Conclusion

With the widescale adoption of hybrid-cloud environments, organizations are asking for a unified solution that help proactively maintain business uptime and migrate workloads across private, public and hybrid clouds. Accenture and Veritas can answer these needs with Resiliency Platform by providing a solution to:

- Orchestrate complex recoveries with predictability.
- Ensure DR readiness with nondisruptive testing.
- Achieve any SLO.
- Seamlessly migrate complex, tiered applications between on-premises and the cloud.
- Extend on-premises resiliency to the cloud with a unified hybrid-cloud strategy.

ABOUT ACCENTURE

Accenture is a leading global professional services company, providing a broad range of services in strategy and consulting, interactive, technology and operations, with digital capabilities across all of these services. We combine unmatched experience and specialized capabilities across more than 40 industries—powered by the world's largest network of Advanced Technology and Intelligent Operations centers. With 509,000 people serving clients in more than 120 countries, Accenture brings continuous innovation to help clients improve their performance and create lasting value across their enterprises. Visit us at www.accenture.com.

ABOUT VERITAS

Veritas Technologies is a global leader in data protection and availability. Over 50,000 enterprises—including 99 of the Fortune 100—rely on us to abstract IT complexity and simplify data management. Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas supports more than 500 data sources and over 150 storage targets, including 60 clouds. Learn more at www.veritas.com. Follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).

2625 Augustine Drive, Santa Clara, CA 95054
+1 (866) 837 4827
www.veritas.com

For specific country offices and contact numbers, please visit our website.
www.veritas.com/company/contact

