DISASTER RECOVERY ASSURANCE

White Paper:
Disaster Recovery Assurance

This white paper introduces the key concept of Disaster Recovery (DR) assurance. Today's DR plans must provide continuous verification and ensure that protection processes remain current and capable of delivering Service Level Agreements (SLAs) to line of business and application owners. In the following paper, we will explore new techniques that give DR processes a level of granularity and automation impossible to attain in physical systems. You will gain insight as to how to deliver Business Continuity / Disaster Recovery (BC/DR) metrics that are meaningful and easily audited, while successfully recovering n-tier applications and IT services in case of disaster or disruption.

This document is intended for:

CIOs, CTOs

IT Directors, Datacenter Managers

Security and Risk Management Officers

BC/DR Planners

UNITRENDS

Play IT Safe

Executive Summary

The spectrum of IT risk scenarios is constantly evolving and challenging CIOs and their staff to comply with corporate business continuity policies, sector-driven standards, and disaster preparedness legislation. Society's ever increasing reliance on IT for its well-being makes it compulsory for organizations of all sizes to demonstrably recover from disruptions and disasters in their IT infrastructure. DR is a board-level concern.

While virtualization and cloud are transforming IT service delivery, the rate of change to infrastructure is accelerating, sometimes out of control. Hardware upgrades, software updates, middleware patches, security and protection mechanisms, and other new components are being introduced at an unprecedented rate, making it all but impossible to gauge the risk exposure of individual changes.

In parallel to the more dynamic infrastructure, n-tier application complexity is increasing geometrically as the number of software components that collaborate in the provision of an IT Service continues to grow. New application architectures fuel interdependency and the risk that a component failure will have unexpected effects on other components, leading to the possibility of severe service disruption.

Whatever the cause of outages, assuring recovery and continuity can only be done through a much higher level of automation and orchestration across the different layers of the service delivery stack, from the network to the application. Legacy mechanisms for DR testing are mostly manual, expensive, disruptive and infrequent—yearly, for most organizations. Unitrends proposes a new paradigm where DR testing becomes application-centric, fully automated and iterative, with daily or hourly cycles that does not interfere with production.

This white paper introduces the concept of DR Assurance in response to the need to continuously verify and ensure that protection processes remain current and capable of delivering Service Level Agreements contracted with the line of business and application owners. We will explore new techniques that give DR processes a level of granularity and automation impossible to attain in physical systems. You will gain insight as to how to deliver BC/DR metrics that are meaningful and easily audited, while successfully recovering n-tier applications and IT services in case of disaster or disruption.

Intended Audience

This document is intended for:

- CIOs , CTOs
- IT Directors, Datacenter Managers
- Security and Risk Management Officers
- BC/DR Planners
- Datacenter Architects
- Application Owners

DR, Virtualization and Hybrid Cloud

Floods, fires, earthquakes, and terrorism have long been top-of-mind disaster scenarios. These large scale situations have brought about a tendency towards an all-hazards, all-inclusive approach to DR.

This approach is so costly and complex that DR budgets can run into the tens of millions of dollars. Yet DR tests are only carried out once or twice per year, with new recovery issues uncovered in every test. The awareness of recovery risk is on the increase, but few organizations have a comprehensive plan to address these concerns.

Currently, sites that can afford to test are fortunate. But for organizations who back up their systems but cannot afford to test DR preparedness, a real disaster situation can be a harrowing experience. At best, restore processes fail to complete automatically and require manual intervention, and at worst cause severe business disruption and significant exposure to contractual penalties or legal action.

Virtualization and cloud enable new paradigms and approaches to DR. Virtualized workloads are independent of the physical hardware, and certain disruptions can be handled automatically or very quickly. For instance a memory failure in a server may not generate a disruption if the hypervisor provides fault tolerance, and transfers the workload to the next available processor without loss of service.

Because virtual assets are streams of bytes, managing them is similar to managing very large files. This similarity paves the way for the introduction of innovative DR practices where the scope of disruptions that can be handled automatically is significantly larger than for physical data centers.

Virtualization enables the atomic unit of recovery to be a virtual machine (VM), where n-tier applications run across several VMs working in sync.

Figure 1 (following page) depicts two typical n-tier applications (IT services).

Each box represents a VM. Applications are run from web servers, application servers and an ERP backend. Each IT service is delivered through a collection of VMs, and two VMs are shared across both IT services.

Figure 1: IT Services and components

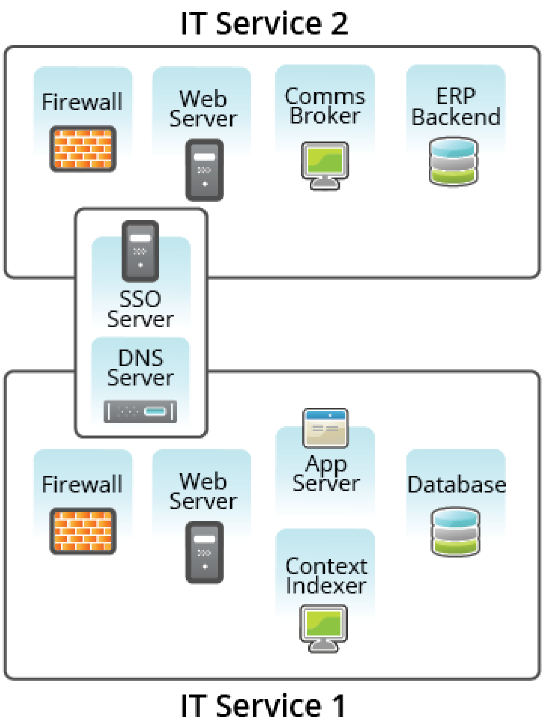

In a virtualized environment, the precise location of VMs can change as hardware is constantly repurposed and shared across a large number of IT services. Virtualization enables better utilization of hardware and environmental resources, but also stretches legacy DR software and processes beyond what they were designed to do.

Challenges in Assuring Recovery

With the advent of composite and n-tier application models, SOA and new delivery models like SaaS, traditional data-based restore methods are becoming imprecise and in some cases obsolete, for the following three reasons:

1. Inconsistent Recovery: when individual components in an IT service are backed up independently, or replication ends abruptly during an incident, there is significant risk that the restore process will bring up components out of sync. In this case, applications will behave inconsistently or unpredictably, or even fail to start, requiring extensive manual intervention to bring them back into production.

2. Fragile processes: huge increases in data volumes coupled with 24x7 IT service operations cause backups to be often split or adjusted to accommodate the backup window. When backup jobs fail to complete, there is rarely enough time to restart them, and therefore the state of the last available point-in-time of n-tier applications is unknown.

3. Disaster propagation: backup and replication copy data from one device to another and is totally service-unaware. Therefore, disasters of a logical nature such as database corruption, bad patches or malware get copied to the DR site, where their presence can go undetected until a test is run, or a disaster occurs.

In order to ensure service recoverability the above problems must be tackled:

1.Component interdependence: components of applications and IT services must be safeguarded in a consistent state that is guaranteed and recoverable according to recovery policies.

2. Process errors: DR must have a closed-loop mechanism that supervises the successful completion, and initiates immediate remediation in case of failure.

3. Disaster containment: must be discarded immediately and replaced with healthy checkpoints.

Technology exists today that provides those solutions for virtualized applications.

Replication and Virtualization

Snapshot and replication have been used by large enterprises for a decade or longer, and in recent years have rapidly penetrated mid-size companies and SMBs due to significant price reductions. When applied to virtualized environments, arrays can snapshot and replicate VMs on the fly with short Recovery Point Objectives (RPOs).

However, storage array replication typically requires similar hardware in both the primary and the secondary sites. Use of dissimilar hardware is possible through storage virtualization or host-based (software) replication (HBR) across sites, as VMs are flat files. Additionally, many backup solutions, including Unitrends, also include replication of backup data that can be used to generate fully hydrated replicas at a secondary site. This variety of replication techniques gives datacenter architects more choices to comply with RPOs and RTOs, particularly if there are budgetary constraints precluding the use of similar hardware at all sites.

While HBR and backup replication are less efficient than array-based replication, advances in technology such as Instant Recovery provide acceptable performance for many applications that use dissimilar hardware in production and DR sites. This is giving more choices for datacenter architects to use lower-cost hardware at recovery sites, lowering DR costs with acceptable RTOs.

Recovery Compliance

The beginning of the twenty-first century has witnessed large-scale disruptive events that have underscored modern society's reliance on essential public and private services. The protection of these services is an area of broad legislative activity where governments are taking an active role in enforcing the deployment of Business Continuity Management (BCM), like the Civil Contingencies Act in the UK. Government bodies like the European Central Bank and standards-setting organizations are also actively issuing guidelines or preparing BCM programs, such as BS25999 or PS-Prep or ISO/IEC27001.

In addition to complying with laws, regulations and standards, CIOs have to consider how to facilitate e-discovery to address litigation, as more and more events and transactions are recorded exclusively in electronic form.

Virtualization enables a new paradigm to address these concerns. VMs can be snapshot simultaneously representing an IT service frozen in time, digitally "cryo-stored". This IT Service can be restored at a later point in time, and executable from that point forward, regardless of the underlying hardware where they are stored or executed.

This enables data center managers to provide long-term storage for an IT service, data and applications, with the assurance that they can be restarted on different hardware anytime in the future. Compliance testing and e-discovery can be quickly and inexpensively facilitated in a virtual or cloud environment.

Figure 2: ReliableDR Multi-site Orchestration

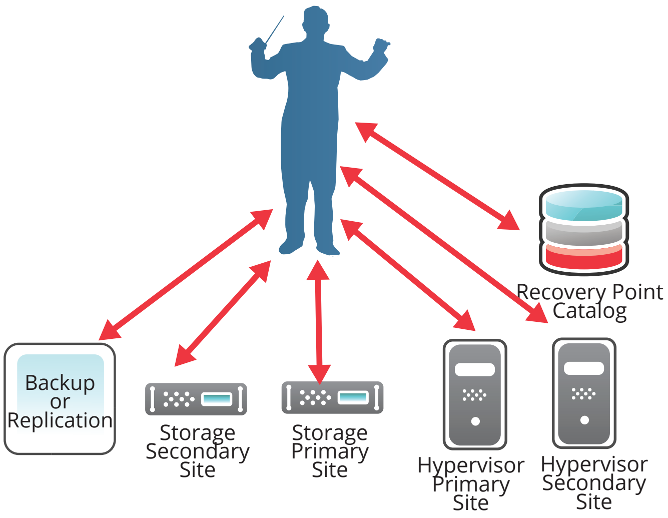

In this environment it is also possible to add recovery compliance processes, such as user acceptance tests or RTO certification, after the verification that a replicated service is in a consistent state. Compliant application checkpoints can therefore be kept on stand-by in off-premise data centers or clouds in case that fast failover is needed due to the loss of a production data center.

Increasing IT Service Resiliency

Each of the enabling technologies described in the Replication and Virtualization section can be programmatically driven to introduce DR orchestration, creating innovative and 100% automated processes that increase IT service resiliency.

DR orchestration brings together the virtualization layers in hypervisors and storage arrays, along with replication or backup software, to manage consistent and recoverable copies of IT services.

Preparing a DR Test

Sets of VMs can be snapshot or restored simultaneously, but are not necessarily in a consistent, recoverable state. As an example, a snapshot could occur in the middle of a transaction where the database considers the transaction complete, but the web server front end still has it in memory. Therefore, snapshots are candidate recovery points until proven that the entire set of VMs is a consistent state.

To test VM snapshots, DR orchestration can leverage the programmability of the hypervisor layer, creating a software-defined datacenter (SDDC) using VMs. Snapshots can then be moved and the IT service run to simulate a failover scenario using runbook automation. This environment is called a sandbox or a test SDDC.

Depending on how the service validation is carried out, a test sandbox can be configured either exactly like the primary site, reusing IP addresses within a closed networking environment, or with different IP addresses enabling some network traffic in and out of the sandbox. It is also possible to include some service VMs that can replace physical resources during a DR test.

From a dynamically built sandbox, DR orchestration can validate the candidate recovery point. This is achieved by starting the VMs in the sandbox using control flow policy to protect IT service consistency. The heartbeats of the VMs will be checked and individual Windows services and Linux daemons will be verified to be running. This process checks the health of each component VM recovered from snapshot.

Recovery Runbooks

DR orchestration can drive the validation of candidate recovery points further through runbooks that include the execution of commands, queries, transactions or any other relevant workload that proves the operational readiness of the recovered application(s). Since the sandbox is isolated from production, virtually any test can be run against the recovery point candidate, from the trivial to the most thorough tests.

DR testing becomes a standard part of the corporate compliance process managed by a specialized team that develops specific recovery runbooks to certify the recoverability of each IT Service. Recovery runbooks validate each tier of an application and apply business rules to each recovery step ensuring policy compliance. As an example, recovery runbooks can be created to perform component-level validation first, and extended to perform in-depth IT Service unit testing at a later date.

Figure 3: Legacy DR Process

Continuous Dependencies increase resilience and speeds of DR testing and recovery because failovers start from a common baseline, certified for multiple applications that have components in common.

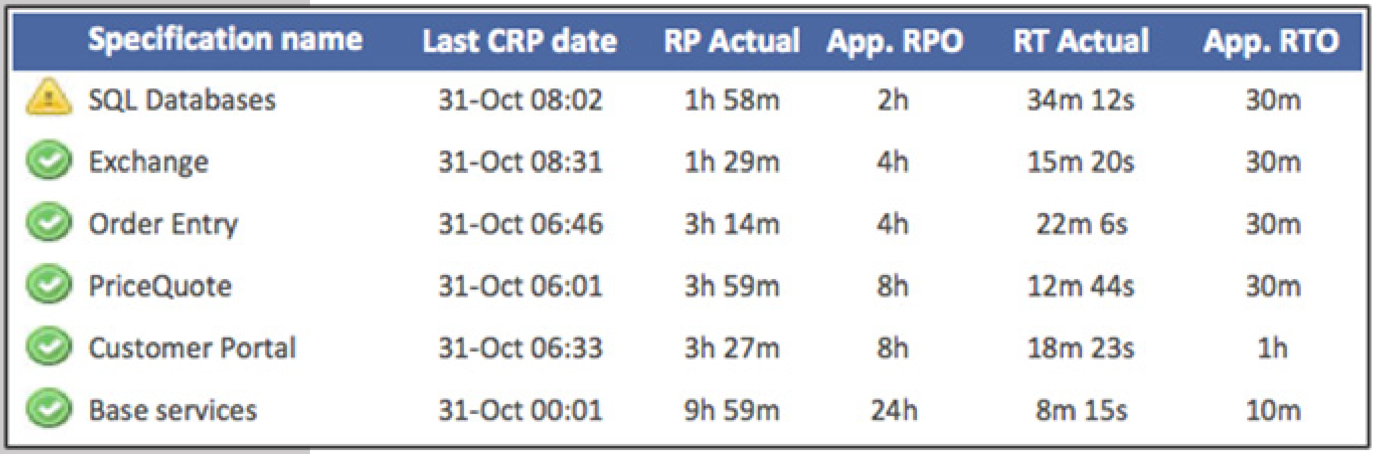

	Specification name	Last CRP date	RP Actual	App. RPO	RT Actual	App. RTO
⚠	SQL Databases	31-Oct 08:02	1h 58m	2h	34m 12s	30m
✓	Exchange	31-Oct 08:31	1h 29m	4h	15m 20s	30m
✓	Order Entry	31-Oct 06:46	3h 14m	4h	22m 6s	30m
✓	PriceQuote	31-Oct 06:01	3h 59m	8h	12m 44s	30m
✓	Customer Portal	31-Oct 06:33	3h 27m	8h	18m 23s	1h
✓	Base services	31-Oct 00:01	9h 59m	24h	8m 15s	10m

The degree of recovery testing within a sandbox is only limited by the amount of virtual resources available, and the period of time available for tests. There is no danger of contamination on the production site and when the tests are complete, the sandbox can be shut down and discarded.

After running the recovery policy and verifying completion, the candidate recovery point can be certified and kept. The DR orchestrator timestamps and makes a catalog entry for the set of snapshots as a single certified recovery point, ready for use. Furthermore, Recovery Time Actuals (RTAs) can be compared to RTOs, to ensure compliance with SLAs or contractual commitments.

Closed-Loop DR

In the event of a failed test the candidate recovery point is invalid and the DR orchestrator closes the loop by discarding the snapshots, logging the test failure, and creating a trouble ticket (SNMP trap and/or email) so that an operations analyst can be called to examine the logs and diagnose the reason for the recovery failure.

The dynamic nature of IT services, where changes are being made constantly to hardware, operating systems, middleware and applications, requires this closed-loop mechanism in order to assure recovery. Recovery Point Objectives (RPOs) are just that, objectives, unless the loop is closed, when objective become guarantees and candidate recovery points become certified recovery points (CRPs).

Continuous Dependencies

It is common for different IT services to share several components, such as databases, identity and access managers, middleware servers etc. To assure consistency in the recovery, all IT services are managed by the DR orchestrator as one set, snapshotting all VMs at the same time and testing the recovery together in the sandbox.

The side effect is that the RPO for the entire set of IT Services has to be the same, causing the amount of virtual resources needed for the sandbox to become large. As an alternative, there are cases where it is preferable to allow each service to be orchestrated separately, using a common CRP for the shared components.

Components	**t_0**	**t_0+2**	**t_0+6**	**t_0+8**
Database	snapshot	snapshot	snapshot	snapshot
Application A	snapshot	snapshot	snapshot	snapshot
Application B	snapshot	X	X	snapshot

Continuous Dependencies is a technique where the DR orchestrator loads an existing (baseline) CRP and maintains it available for other components to be tested with it. For instance, there may be two applications sharing a database. Application A has an RPO of 2 hours, and application B has an RPO of 8 hours. In this example, the database and application A are DR tested every two hours. Every 8 hours, the DR orchestrator would snapshot application B separately, and load the available CRP from the last test of application A to certify recovery using the same database.

Figure 3: Supporting different RPOs for services that have components in common

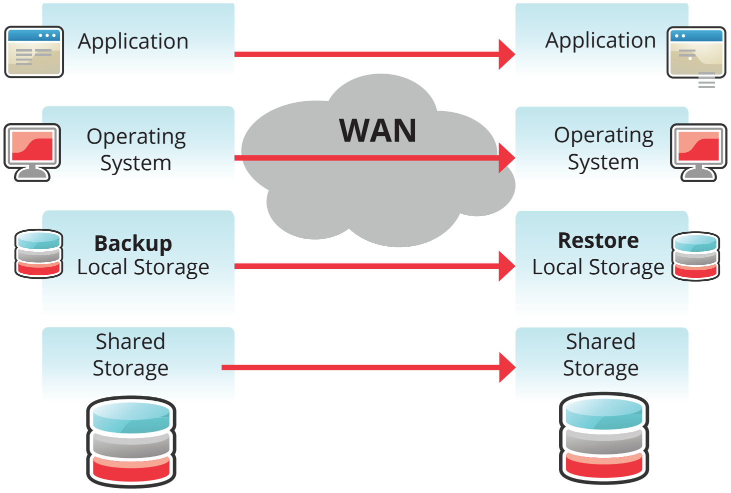

ReliableDR: Modernizing DR

ReliableDR is the next-generation, closed-loop orchestrator for virtual environments that delivers disaster recovery assurance. ReliableDR leverages virtualization, replication, backup and intelligent storage arrays for automating recovery assurance processes that set RPO by business policy, and reduce RTO to minutes while tracking against objectives with real tests.

The core benefit of ReliableDR is its ability to align corporate Business Continuity Policy (BCP) and DR. ReliableDR simulates failover scenarios as often as required, certifies that recovery points are ready for use in case of disaster or major disruption, and automates the failover processes to the point where service is restored.

Legacy DR processes were built under the premise that the infrastructure had to be mapped 1:1 from the primary to secondary sites, and the storage was backed up or replicated periodically. Recovery was slow and labor-intensive. Hardware had to be verified, and operating systems had to be booted individually.

Written DR procedures, when they were available, are often out of date. RTO was constantly on the increase due to:

- **Configuration Drift:** this is the unavoidable result of the very large rates of change at data centers; 1:1 mapping cannot be maintained and configuration differences between production and DR sites need to be calibrated after each DR test, or during a live failover by subject matter experts.
- **Multi-Stage Recovery**: full tests require separate, labor-intensive, and lengthy steps to start the hardware, operating systems, and the applications; separate skills are normally needed for each stage of recovery.
- **Intricate Software Dependencies**: startup procedures need to follow specific sequences so applications and services are started in the correct order and operational before others are launched.

Stopping Configuration Drift

ReliableDR provides configuration-aware orchestration. It has the intelligence to configure recovery VMs in real time, replicating the exact configuration of VMs in the production system at the time of the last certified recovery point (CRP). If the primary site's VMs change, e.g. if memory is increased in one of them, the change is reflected immediately.

With ReliableDR, configuration drift between primary and secondary sites is detected upon execution of the next testing cycle. RPO policy is enforced, and a trouble ticket is raised if any configuration anomaly is detected.

Application-centric, Automated Recovery

ReliableDR orchestrates storage hardware and hypervisor components to produce CRPs on the secondary site. It allows each application to be certified for recovery in a controlled and configured sequence.

Snapshots generated by ReliableDR are consistent and complete, and contain all the necessary components required to restore the entire IT service, including data, operating systems, middleware, web front ends, etc.

ReliableDR allows DR planners to apply business policy to the DR processes of IT services individually, depending on criticality, laws, regulations or board directives. For critical IT services, business polices can be applied every 30 minutes to 1 hour if desired, and ReliableDR will generate 24 or 48 CRPs per day.

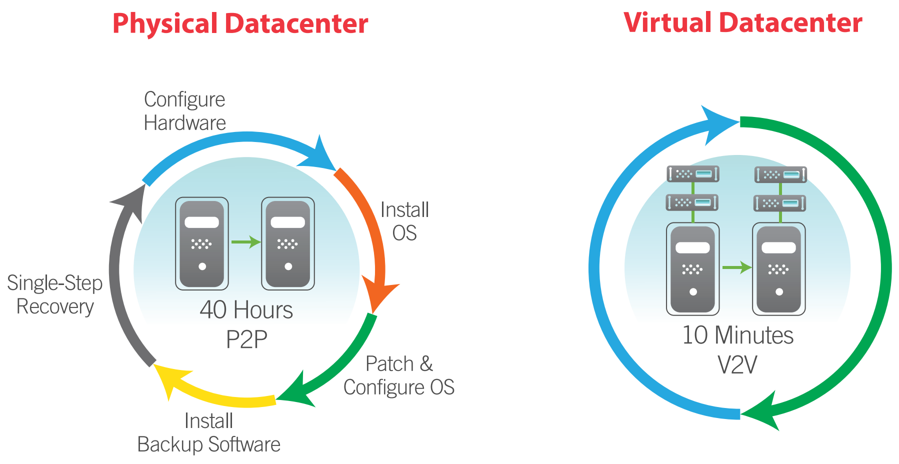

Figure 5: Comparing Legacy and Virtual DR

	Legacy	Virtual
Test Frequency	Yearly	Daily
Granularity	Files	IT Services
Mechanism	Backup	Replication
RPO	24 Hours	<1 Hour
RTO	Days	Minutes

Because recovery is fast and automated, data centers only need enough hypervisor licenses in the secondary site to certify the individual recovery points. In case of a real contingency, the licenses used in the production datacenter can be used for recovery.

RTO Acceleration

ReliableDR can configure, mount, and start certified recovery points CRPs very quickly, usually in minutes. ReliableDR enables multiple IT Services to be restarted in parallel while respecting service interdependencies through flow control logic that enforces boot order priorities. This added intelligence reduces RTO, requires fewer skills, and less documentation, intervention and training on the part of system administrators in the case of a real contingency.

The improvements in RTO can be dramatic. Coupled with the ability to policy-drive the RPO generation, DR orchestration brings close alignment between IT disaster recovery processes and business policy.

Post-Certification Processes

ReliableDR provides support for additional, customized processes to be executed using Microsoft PowerShell and Linux bash scripts. This capability can provide many advantages, including:

1. **Move vaulting from the primary to the secondary site:** for data centers where long term backup retention is mandatory, the process can be moved to the recovery site and effected on certified recovery points. This not only reduces overhead on the primary site, but gives auditors and investigators the ability to restore fully operational points-in-time that include data and applications.
2. **Additional compliance processing:** organizations with more intricate regulatory or statutory compliance requirements can run additional processes and extract recovery reports on the certified recovery point CRP as needed without impacting production services.
3. **Business analysis:** LOBs can mine live data and run reports at any time without affecting production performance.
4. **Test and development:** sandboxed replicas of complete and fully working IT services can be made available quickly for test and development without affecting production.

Failover & Failback

At the time of a real contingency, datacenter management will instruct ReliableDR to switch from testing mode to live failover mode. Instead of using a sandbox, ReliableDR generates VMs with normal network connectivity.

Using the catalog of CRPs, systems administrators choose the snapshot to recover from, and ReliableDR proceeds to start up the snapshot VMs using the control flow logic to restore the IT service in the correct sequence needed.

ReliableDR guarantees that there is always a recovery point that is consistent and fully recoverable, and ready to provide service continuity.

When the production data center is available again, the failback process can start. ReliableDR will assist or automate the reversing of the mirror and start to test the recovery process in the production site. Upon verification that the failback procedure is working accurately, the secondary site will be shut down, a last replication effected, and production will resume at the original data center.

Case Studies

Santalucia Insurances

Santalucia is an insurance company with a network of 365 agencies, 7.5 million customers and 9,000 employees. Headquartered in Madrid, Santalucia provides home, life, health and accident insurance, as well as a wide portfolio of financial products for investment, pension and retirement plans.

An early adopter of virtualization, Santalucia has primary (production) and secondary (DR) data centers. Each has eight x86 farms hosting critical applications; storage is centralized in each datacenter using HP XP arrays. Using the HP XP native snapshot and replication capabilities, the production virtual machines are hot-copied regularly to the secondary datacenter.

ReliableDR orchestrates the VMware hypervisors and the HP XP storage arrays. To ensure consistency and immediate availability in case of disaster, ReliableDR verifies that the VMs are configured at the secondary datacenter exactly like the primary site. This is done for over 100 VMs on a daily basis.

For a set of 60 VMs deemed to be mission-critical, ReliableDR runs additional recovery point certification jobs every night between midnight and 4am. These jobs create VDCs in the secondary datacenter and bring up all 60 VMs. Upon successful execution and testing of each IT service, the snapshots are certified and cataloged. Santalucia retains in ReliableDR seven generations of Certified Recovery Points from 24 hrs to 7 days that can be executed in the event of a disaster.

DR orchestration takes place out of the secondary datacenter, where ReliableDR runs as a virtual machine. The production datacenter is not impacted in any way by the DR orchestration. In case of a disaster, be it the entire primary datacenter or a subset, IT services can be brought up selectively and automatically by ReliableDR.

Van Lanschot Bank

Van Lanschot, founded in 1742 and one of the oldest private banks in Europe, is an independent financial institution that offers private banking, asset management, business banking and corporate finance. It is headquartered in the Netherlands and has operations in Belgium, Luxembourg, Switzerland and Curaçao.

Because the bank operates across several jurisdictions, it faced the challenge of demonstrating legal compliance from several countries, such as the UK Financial Services and Markets Act (FSMA), as well as being subject to regulations and audits from the National Bank of Belgium (NBB) and the Dutch National Bank (DNB).

Van Lanschot

The bank runs two NetApp Filers in its production and DR data centers and was using SnapMirror in order to generate recovery points. However, the existing toolset did not provide fully automated and iterative DR testing, which the bank needed in order to demonstrate compliance at all times.

In particular, there was a need to reduce the existing RPO below 24 hours. A search was made for advanced tools that could leverage their recent investment in a new virtualized infrastructure.

ReliableDR was deployed at Van Lanschot's DR site and integrated with SnapMirror and FlexClone to automate DR testing. Several RPO tiers were defined, and for mission-critical applications ReliableDR runs twelve DR tests per day and provides an iron-clad RPO of two hours. Out-of-the-box recovery certification functionality was applied to Microsoft SQL Server databases and Exchange. Corporate application servers were certified for recovery using ad hoc application tests as specified by the line of business.

Van Lanschot now benefits from tiered and fully automated DR testing of their x86 applications portfolio. SLAs are constantly tracked and enforced. Paul Timmermans, Managing Director at Van Lanschot Belgium, said, "Our Disaster Recovery strategy cannot fail to deliver and this is why we chose ReliableDR. It lets us run non-disruptive, frequent, scheduled Disaster Recovery tests to ensure that a successful service recovery is always assured."

Conclusion

ReliableDR is a unique solution that eliminates DR testing processes and guarantees that RPO and RTO for IT Services and business applications can always be met, whatever the scenario.

Whether improving disaster recovery capabilities, reducing cost, demonstrating compliance or moving towards a hybrid cloud model providing central resiliency to distributed environments, ReliableDR enables data centers to:

- Align business continuity policy with IT disaster recovery.
- Deploy fully automated, non-disruptive and continuous DR testing.
- Guarantee that all components of an IT service are always fully recoverable.
- Report compliance deviations via dashboard, email or problem management systems.
- Measure and enforce RPO and RTO automatically.
- Failover individual applications at the push of a button.
- Scale recovery from individual applications to large data centers.

About Unitrends

Unitrends provides physical, virtual and cloud-based protection and recovery for every organization's most valuable assets: its data and applications. Supported by a "crazy-committed" customer service model based on engagement, experience and excellence, the company consistently achieves a 98 percent customer satisfaction rating and lets everyone play IT safe by delivering the best cost-to-value ratio in the data protection and disaster recovery industry. Visit www.unitrends.com.

Unitrends
7 Technology Circle, Suite 100
Columbia, SC 29203
1.866.359.5411
sales@unitrends.com

Try Unitrends ReliableDR FREE. Download the trial at:
www.unitrends.com/products/software/reliabledr