UNITRENDS

Cloud-Based Recovery Assurance

IS YOUR CONFIDENCE IN YOUR ABILITY TO RECOVER CRITICAL DATA AND APPLICATIONS SHRINKING?

Cloud Based Recovery Assurance

Is your confidence in your ability to recover critical data and applications shrinking?

End users have come to see IT as a service, like water or electricity. They expect every company's data and applications to be available at all times, from any location and any device. Customers and employees don't even think about it until it's no longer available. A customer trying to make a purchase through a company's website expects the transaction to go smoothly, whenever they're ready to buy and from any device they happen to be using.

As an IT professional, this is hardly news. But what is surprising is the gap that exists between these expectations and the reality within many IT organizations. While IT professionals are doing all they can to meet these growing needs and challenges, an integral component of their mission—backup and disaster recovery—fails to receive the strategic focus required. While many organizations have some level of disaster recovery (DR) plan in place, due to cost and other factors they aren't testing them as often as needed to assure recovery when an outage or catastrophic failure occurs.

When that outage or disaster happens and you need to quickly restore your company's data and IT services, an untested DR plan can pose a serious threat to the company as a whole. And for you, it can mean the difference between being seen as a hero to the C-suite, or updating your résumé in search of another job.

A local government agency in New Orleans offers a real world example. Given their location, it's obvious they are prone to natural disasters, so a DR plan was in place. However, when two servers that held the parish's conveyance and mortgage records dating back to the 1980s crashed simultaneously, their DR plan came up short.

The Times-Picayune newspaper[1] reported on the incident, saying, "Without a complete and verified database of both conveyance and mortgage records, title companies can't be sure that a person trying to sell a property truly owns it free and clear. And the mortgage record database, which is separate from the one for conveyance records, is still missing about 100,000 documents."

The parish's IT staff had thought it was backing up the parish's data and replicating the data to a cloud backup and disaster-recovery-as-a-service (DRaaS). But neither the IT staff nor the service provider was aware that the data hadn't been backed up for months. What data that had been backed up had passed its 30-day retention policy and was deleted. Without proper testing of the data and applications at both the primary data center and in the cloud, there was no validation of backup or assurance of recovery of the IT services, even from a garden-variety hardware failure. What happened in the agency in New Orleans illustrates a situation that is all too common and costly but yet completely avoidable. With the right tools you can quickly implement backup recovery assurance and have the peace of mind that your business will continue in the wake of events that would otherwise be catastrophic.

In a Forrester Research report titled "Develop Your Road Map for Business Technology Resiliency Tools,"[2] Forrester reports that "due to poor planning and testing, the majority of companies aren't prepared for a disaster." For a range of reasons, many companies have not made the proper investment in mitigating their most common causes of downtime, which continue to be mundane events such as power failures, IT hardware failures, network failures, and human error." Today this effect is compounded by the increasing demands for 24x7x365 availability to critical applications.

As an IT professional, you will experience a system failure, outage, or complete site disaster. It's inevitable. Your organization may already have a full or partial DR plan in place, but are you certain you can recover your critical applications and data in a timely manner that meets corporate business continuity requirements and customer expectations?

[1] "Real estate computer crash brings industry to its knees" Dec 5, 2010 http://www.nola.com/politics/index.ssf/2010/12/real_estate_computer_crash_bri.html)

[2] "Develop Your Road Map For Business Technology Resiliency Tools" Rachel Dines, July 24, 2012 https://www.forrester.com/Develop+Your+Road+Map+For+Business+Technology+Resiliency+Tools/fulltext/-/E-RES57333?objectid=RES57333

Understanding Downtime and Availability

Most organizations define "availability" somewhere along a continuum between multiple hours of downtime with significant data loss to real-time 24/7 uptime with zero data loss. Your definition depends on your business needs, data and application requirements, and organizational structure. The goal, however, should be to prevent the inevitable system downtime from affecting critical business uptime. There are two types of downtime: unplanned and planned.

Unplanned Downtime

Surprisingly, unplanned downtime represents less than 5-10 percent of all downtime. These events include security violations, corruption of data, power outages, human error, failed upgrades, and natural disasters.

Unplanned downtime can strike at any time from any number of causes. Although natural disasters may appear to be the most devastating cause of IT outages, application problems are the most frequent threat to IT uptime. According to Gartner, people, hardware failure, and process problems cause an estimated 80 percent of unexpected application downtime. Human error, such as not performing a required task, performing a task incorrectly (software configuration errors), over provisioning storage, or deleting a critical file, can play havoc with applications.

Planned Downtime

While unplanned events tend to attract the most attention, planned downtime actually poses a bigger challenge to business uptime. Routine daily/weekly maintenance to databases, applications, or systems can lead to interruption of services. Studies show that system upgrades, performance tuning, and batch jobs create more than 70-90 percent of the average company's downtime.

Although companies must be concerned about natural disasters, the inherent daily threat posed by application problems and human error should be a major focus when planning for inevitable downtime. This is especially true when the exposure of software applications to unplanned downtime is aggravated by requirements to retain, protect, and audit email, financial, and other records under a growing list of compliance mandates. Add to the mix increased threats of malicious cyber-attacks, the growth of more complex n-tier applications spread across multiple platforms and operating systems, and continued reductions in IT staffing and it's easy to see why business continuity planning is more critical than ever.

Over the past 24 months, 91% of IT organizations experienced unplanned downtime at an average cost of $7,900 per minute. Are you prepared to mitigate these outages and eliminate the devastating cost to your business's revenue and profitability?

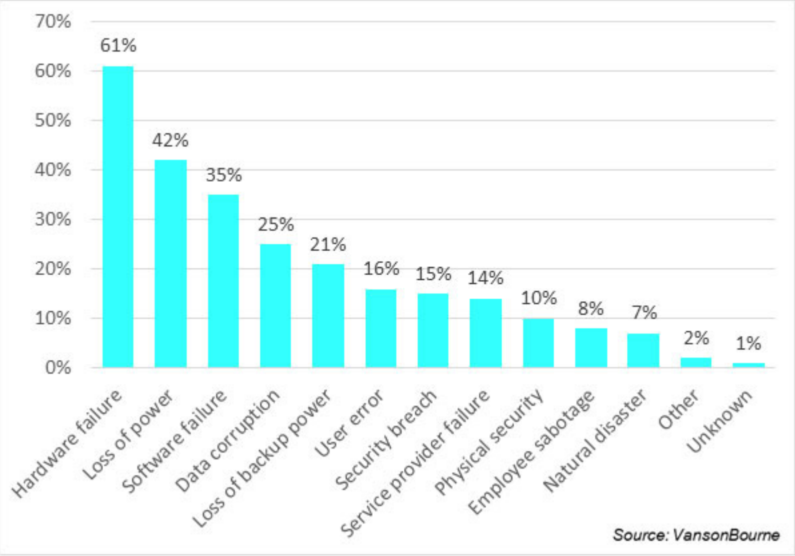

Source: VansonBourne

Figure 1: Causes of System Downtime and Data Loss

Financial Impact: Cost of Downtime

To assess the impact of downtime, it's imperative to understand the effect of downtime cost on the business. Unexpected IT outages can unleash a series of direct and indirect consequences both short term and far reaching. The dollar amount that is assigned to each hour of downtime varies widely depending upon the nature of the business, the size of the company, and the criticality of the failed IT system to primary revenue generating processes.

The following chart includes both direct productivity and revenue loss as well as indirect cost such as damage to reputation, loss of customer opportunity and damage to the company brand. When calculating loss, each of these factors should be included in downtime cost calculations. An average estimate, according to studies and surveys performed by numerous IT analyst firms, downtime cost businesses between $90,000 and $300,000 (US) for every hour of IT system downtime. Loss to large financial institutions, telecommunications, manufacturing and energy companies can be significantly higher.

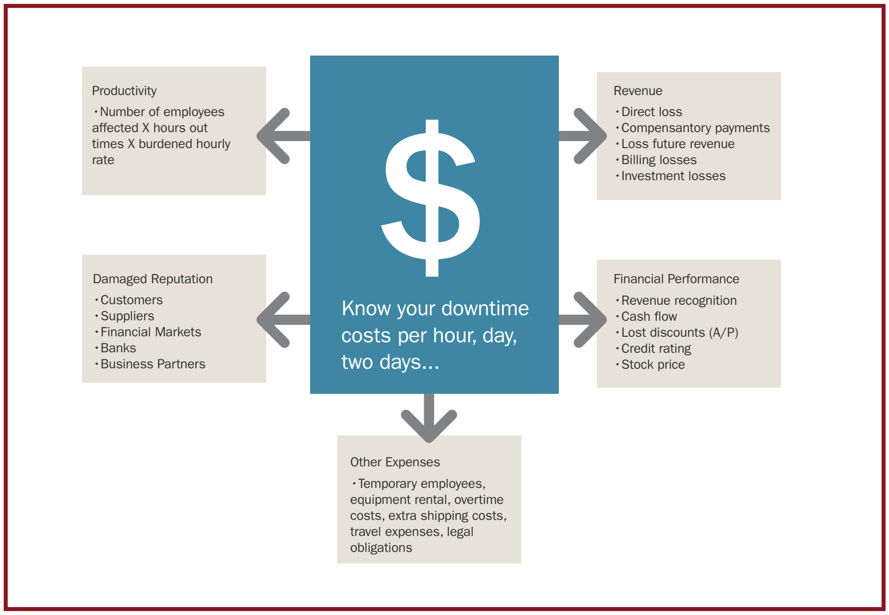

Figure 2: Know Your Downtime Costs

Beyond loss of revenue, productivity, and reputation, businesses today cannot function without computer access and functionality. FEMA[3] reports that most businesses that suffer catastrophic data loss or an extended IT outage go out of business within two years of the disaster.

No matter what the cause, downtime impacts more than day-to-day interactions. It can impact the integrity of databases as well as the applications that use them. Some businesses can survive some data loss, while others are dependent on electronic data interchange, or are required to archive information to meet stringent audit, regulatory, and compliance requirements. Business that employ a global workforce that collaborate around the clock, or provide eCommerce to make sales and deliver customer service 24/7, also cannot afford data loss. While all organizations feel the impact of IT system downtime, organizations that rely on real-time data dependency will suffer most from the effects of any downtime, both immediately after the data loss and well into the future.

[3] "Protecting Your Business" June 15, 2015 https://www.fema.gov/protecting-your-businesses

The Added Burden of Compliance

Numerous industry regulations require companies to support more stringent availability standards. What's more, many compliance requirements, whether directed at specific industries or a broad cross-section of companies, mandate the protection of business data and system availability. Businesses may incur financial or legal penalties for failing to comply with these data or business availability requirements. Examples include:

- **Health Insurance Portability and Accountability Act (HIPAA)**—ensures that only properly authorized individuals have access to confidential patient health data and provides long-term guidelines to secure confidential information. HIPAA mandates a five day maximum turnaround on requests for information.
- **Sarbanes-Oxley Act of 2002**—stipulates that CEOs and CFOs attest to the truthfulness of financial reports and to the effectiveness of internal financial controls. Sarbanes-Oxley mandates a required timeframe in which to report financial results—each quarter and at year-end. Failure to make these deadlines can result in financial penalties.
- **Gramm-Leach-Bliley Financial Services Modernization Act of 1999**—limits access to non-public information to those with a "need to know" and requires safeguarding of customer financial information. Loss of important data can lead to penalties for the financial institution.

What is Recovery Assurance?

Recovery Assurance is having complete confidence that mission critical applications will recover in the time required to meet an organization's IT service demands, no matter the disaster or outage, whether planned or unplanned. Complete confidence is gained by coordination of n-tier complex applications across multiple platforms and operating systems, scheduled automated testing, and real time alerts when critical applications fall outside recovery time objects.

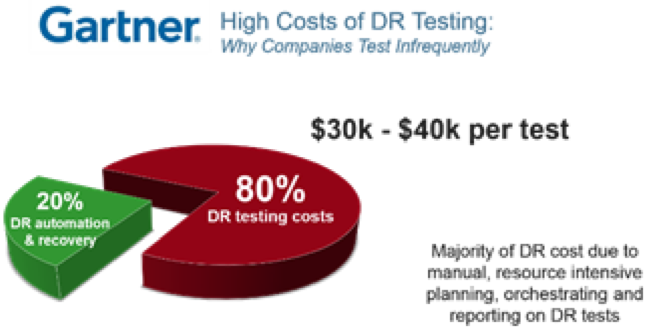

The critical feature of Recovery Assurance is regular and automated testing, a sensitive subject for many IT executives. While few debate the importance of testing, a small number of IT organizations actually take the time to regularly test their plans because of the significant time and cost associated with manual testing. While frequency of testing is critical to recovery assurance and compliance reporting, it's even more important to ensure testing is run across the full software stack of today's complex n-tier application environments. For mission-critical applications, recovery testing must go beyond just testing individual software instances, such as a single database, operating system, or virtualization hypervisor. It's essential to test the interdependency of all instances within crucial applications. This testing includes coordinating data sets, managing boot times, measuring the actual time to recovery (RTA) of the complete mission critical application, and reporting proof of compliance to each specified RTO and RPO (defined in the following section). Following these best practices assures recovery within service levels agreements (SLAs) set by auditors, corporate stakeholders, and customers. Here's the question IT organizations must consider: Following any outage, how quickly must you have the organization up and running as close to normal business operations as possible? Your recovery will depend on two objectives: your recovery time and your recovery point.

Recovery Time Objective (RTO)—RTO defines how quickly you need to restore an application and have it fully functional again. Email and transaction-based applications that are critical to employees and customers will have more immediate recovery times, while applications that are less frequently accessed, such as a human resources applications, may have less immediate recovery time requirements.

Recovery Point Objective (RPO)—RPO defines how much data the business can afford to lose. Applications related directly to business continuity, where data changes significantly every day, will top this list. Back office processes may be lower on the list.

Different applications have different mission criticality for an organization and therefore should have different priorities with regards to RTO and RPO. For example, a supply chain application that feeds a production plant may require a recovery time of a few minutes with very minimal data loss. A payroll system that is updated weekly with only a few records may only require a recovery time of 12 hours and a recovery point of 24 hours or more before the impact will affect the business. Items that can affect recovery time include:

- How many transactions can you afford to lose without significantly impacting revenue and production?
- Do you depend upon one or more mission critical applications such as ERP or CRM software?
- How much revenue will you lose for every hour your critical applications are unavailable?
- What will the productivity costs be for the loss of available IT systems and applications?
- How will collaborative business processes with partners, suppliers and customers be affected by an unexpected IT outage?
- What is the total cost of lost productivity and lost revenue during unplanned downtime?

While defining an RTO and RPO for each mission and business critical application is a good start, solid DR plans must go a step further. In today's explosive data growth environments, IT organizations can no longer depend on static RTO and RPO definitions. Setting and forgetting RTO and RPO targets will negatively impact recovery assurance of mission critical applications in today's stringent corporate business continuity environment. Beyond defining the initial RTO and RPO for each mission and business critical application, it's imperative to regularly test and measure RTOs and RPOs to continuously provide proof that the current DR environment will always meet the business continuity goals associated with the service levels defined by corporate stakeholders and compliance auditors.

Recovery Time Actual (RTA) and Recovery Point Actual (RPA)—The measured time it takes for all application-component interdependencies to boot and become fully available to the business and application users.

Measurement of the RTA begins at the start of a DR test and goes to the moment the last component comes online and users have access to the application. The RTA measurement is used to assure an application is always compliant with the corporate service levels required for that application. As long as the RTA takes less time than the recovery time objective, the application meets business continuity requirements. If the RTA takes longer than the objective, then an immediate alert needs to be sent to the IT team notifying them that an application recovery time objective failed. As a standard part of the automated test process, an in-depth test report should be included with the failure alert so the IT team can quickly determine the cause of failure (reason for being outside the recovery time) allowing the IT team to quickly align the application instance to meet service level agreements.

Recovery Assurance and Your IT Environment

As any good CIO will attest, a DR plan must involve more than just recovering data on servers. Instead, it's about knowing that you can recover your applications with minimal data loss (RPO) and knowing that the applications are able to be online and service clients within a specified time (RTO).

Rather than having to implement separate, siloed backup products for each application at dedicated data centers, applications have evolved to n-tier, complex, service-oriented tools for end customers who expect them to be always-on and accessible from anywhere. The cloud provides an infrastructure to run these applications, as well as accessibility and reduction in the time to deploy while offering virtually infinite amount of scalability.

Additionally, the cloud provides organizations an agile and elastic DR environment where data and applications can be spun up in the event of an outage. With automated testing and coordination of resources, recovery is predictable and free of operator effort. The ease and flexibility associated with recovery assurance in the cloud ensures recovery of critical applications with minimal data loss while maximizing uptime, meeting corporate business continuity requirements.

Recovery Assurance provides stakeholders confidence and peace of mind that backups will no longer disrupt their day-to-day business processes and, when they need data recovered, that they get it back quickly and non-disruptively. But, ultimately, it builds confidence that in the event of an outage the critical aspects of the business will continue and this impacts the bottom line.

Cloud Recovery Assurance requires less hardware and software than existing data center recovery and take less time to manage. This translates into lower upfront capital expenses as well as reduced ongoing operating expenses. Yet the biggest advantage that Cloud Recovery Assurance provides over traditional data center recovery is that it positions the organization to grow, compete, and win. In today's competitive business environment, the ability to easily and quickly recover applications—not just back them up—gives a company a competitive advantage.

Today's IT staff as well as application owners can rightly and confidently expect to recover applications and their data in a timely manner since they have the opportunity to test and certify recoveries, locally and in the cloud. With Recovery Assurance, organizations can implement those long talked about DR plans with confidence that they have a solution in place that is easy to implement and can be automatically tested to assure recovery in the event of any outage, planned or unplanned. Further, Recovery Assurance facilitates application failover, and recovery of physical servers and virtual machines (VMs), giving any size organization access to a robust DR solution that comes with confidence of business continuity and peace of mind.

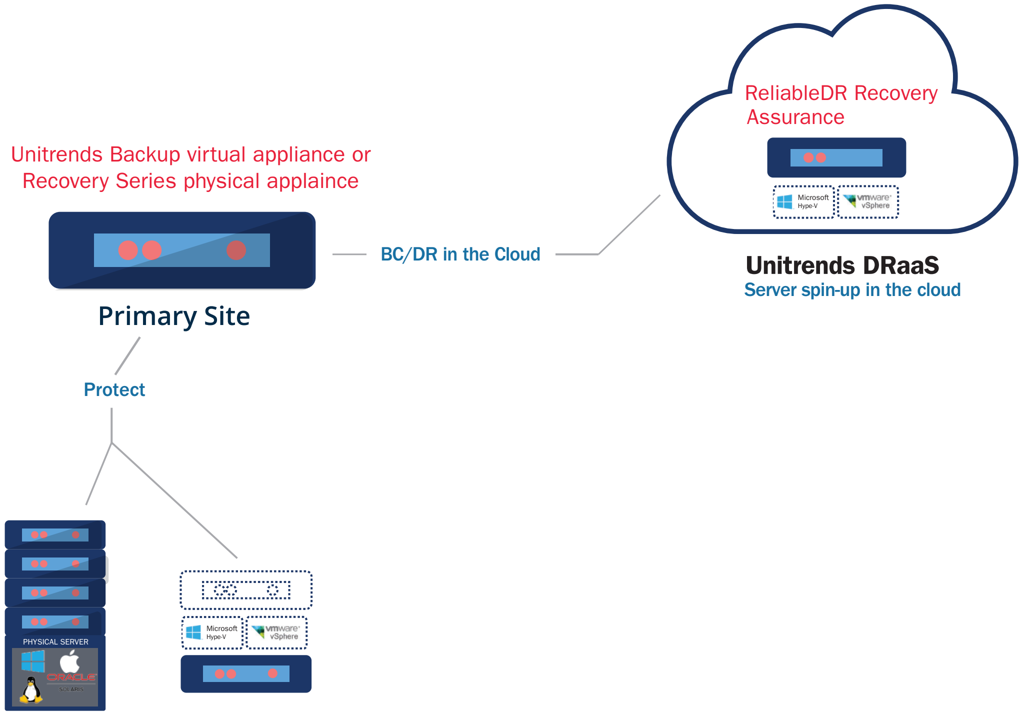

Figure 3: Unitrends Backup and ReliableDR

Recovery Assurance and the Cloud

Historically, DR has been one of the biggest challenges for IT organizations. Massive amounts of time and resources were spent ensuring that the primary and a dedicated remote site remained exactly alike, even in terms of software revisions and patches. Even in best case scenarios DR fell short, and restores failed due to backup failures, system configuration issues, or human error. As a result of this complexity and cost, IT organizations only used secondary site DR to protect applications that were absolutely critical to the survival of the business. Other applications either did not have a DR plan or they were simply backed up on tapes and sent offsite for safekeeping. If a disaster occurred, these tapes were transported back to a secondary site where a new infrastructure was created for recovery. Cost was high but most importantly RTOs and RPOs were measured in days and often weeks. And while the cost of tape was lower than mirroring to a secondary site, it was still too high and there was no guarantee of recovery.

The emergence of cloud-based recovery, virtualization, and improvements in compression, deduplication and WAN optimization have greatly influenced DR best practices, allowing for the new standard of application recovery assurance in the cloud. What used to be a stagnant disaster recovery environment based on a mirrored physical replica is now an agile, elastic extension of an overall cloud strategy. This elastic cloud disruption combined with recovery assurance provides an environment that matches demands for 24/7 IT services, regulatory compliance, and lower cost. This allows IT organizations to focus on protecting workloads and applications with granular RPOs and RTOs and therefore mitigate or eliminate downtime and data loss.

However, even in a dynamic cloud environment, some old rules still apply. Increasing cost of downtime, a growing reliance on IT services, and more aggressive RTO and RPOs make it crucial to understand the criticality of different application workloads, how these workloads change over time, and the impact of data growth on time to data recovery. To guarantee business continuity, best practices combine validation testing with instant recovery of applications in the local data center (on premise) with assured recovery and spin-up of mission critical applications in the cloud.

To support this elastic, growing, 24/7 IT environment, the past “set it and forget it” testing philosophy no longer applies. Like the example of the government agency in New Orleans, the worst time to learn that backups were not completed, data sets changed, or recovery times no longer compliant with objectives is during an outage. To meet the IT service, compliance, and regulatory audit demands, it’s imperative to have an aggressive disaster avoidance testing regimen that coordinates complex n-tier application instances, automates and schedules testing, and measures actual recovery times and recovery points for each backup and critical application. Only this level of automation will allow the IT staff to continuously root out and eliminate potential disasters before they strike. Additionally, cloud-based recovery assurance provides the flexibility to easily reconfigure DR testing routines to match changing application RTO/RPO requirements and corporate service level policies while seamlessly adding capacity to accommodate the evolving business. As business requirements shift, the local on premise and cloud-based DR strategy must have the agility to rebalance and grow with the business.

Disaster recovery is about minimizing downtime and assuring business continuity, providing employees and customers access to critical IT services within defined service levels while also mitigating revenue and data loss. The cloud has become an increasingly important part of the infrastructure for delivering DR agility and scale, allowing IT organizations to protect more workloads and simultaneously deliver reliable and resilient protection and recovery.

Recovery assurance in the cloud adds a higher level of guarantee, not only recovering files and databases, but coordinating and automatically testing heterogeneous applications as well as measuring recovery times and providing compliance proof to auditors and corporate stakeholders. In the event of a disaster or a service outage, the IT staff can recover an entire service instead of just reassembling parts. Recovery Assurance is about bringing an entire IT Infrastructure back online, or at least the most crucial applications and systems that are key to running the business.

Unitrends Recovery Assurance—ReliableDR

In this paper, we’ve covered the importance of automated DR testing and recovery assurance within today’s technology landscape. But the IT marketplace still offers few recovery assurance tools that integrate well with backup and cloud solutions. With ReliableDR Recovery Assurance, Unitrends proposes a new paradigm where DR testing becomes application-centric, fully automated and iterative.

Featuring seamless integration with Unitrends Backup software and Recovery Series appliances, ReliableDR, delivers the industry’s first recovery assurance for physical Windows backups in addition to VMware vSphere and Microsoft Hyper-V applications. IT administrators can now validate the recovery of n-tier applications across mixed virtual and physical environments—a vital need for non-virtualized, mission-critical physical Windows servers.

With ReliableDR, for the first time companies of all sizes have unified, affordable recovery assurance that automates and assures recovery processes across a heterogeneous IT environment reducing risk, decreasing recovery times, and sustaining the business through any outage, planned or unplanned. Assurance that the recovery will be successful and real-time visibility into meeting recovery time and recovery point objectives.

Real Time Management and Monitoring Dashboards

ReliableDR is a comprehensive, application-level recovery assurance framework, centrally managed and monitored by a real-time, interactive web-based console. The console contains tools for enterprise- and job-level monitoring and management.

Enterprise-level monitoring enables consolidated monitoring of all jobs from a central location. Monitoring and management at the enterprise-level allows an IT user to view and edit running jobs, job history, all tasks and events, and scheduled tasks. Job-level monitoring and management enables monitoring of individual jobs. At the job level, graphical summaries can be viewed, job options edited, server details and job history viewed, and server settings edited. Figure 4 is an example of a Summary Tab for viewing a mixed virtual and physical service job detail and settings. The ring graphs summarize test history, VM RPO compliance, and VM RTO compliance. Green represents successful tests.

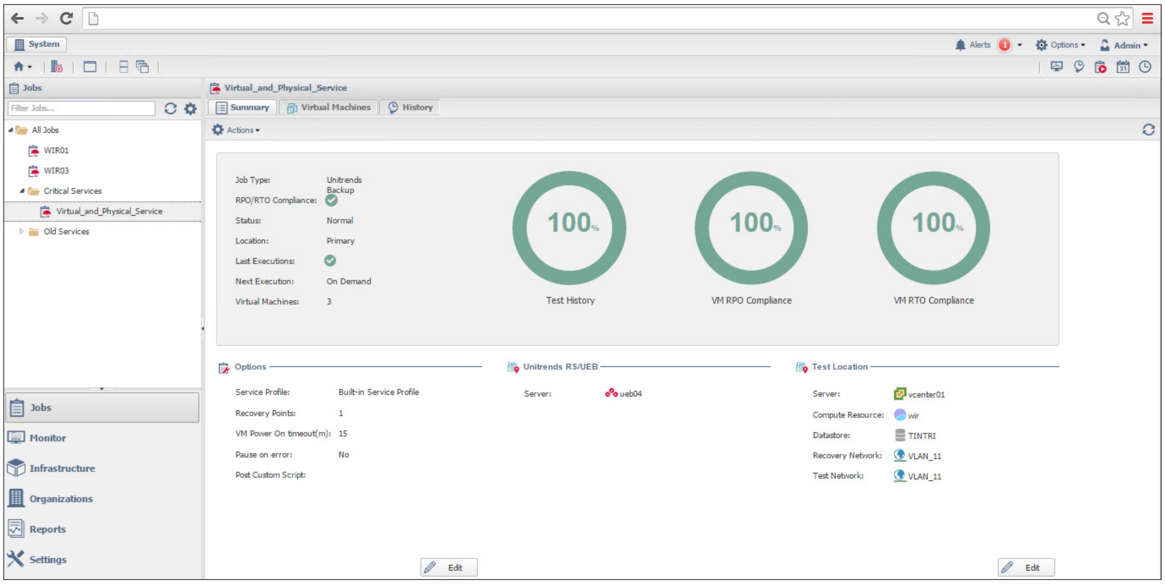

Job Type: Unitrends Backup
RPO/RTO Compliance: ✓
Status: Normal
Location: Primary
Last Executions: ✓
Next Execution: On Demand
Virtual Machines: 3

Options
Service Profile: Built-in Service Profile
Recovery Points: 1
VM Power On timeout(m): 15
Pause on error: No
Post Custom Script:

Unitrends RS/UEB
Server: ueb04

Test Location
Server: vcenter01
Compute Resource: wir
Datastore: TINTRI
Recovery Network: VLAN_11
Test Network: VLAN_11

Figure 4: Top Level Summary Dashboard

Recovery Jobs: Application Recovery Assurance

Recovery isn't successful until a failed critical application is operational again and users are up and running. In many cases, this isn't just one VM, but multiple instances and VMs work together within a network to deliver business services. ReliableDR tests those situations—not just recovery of a single VM, but the entire IT service. ReliableDR performs compliance tests, creates certified snapshots that can be used to quickly recover virtual (VMware and Hyper-V) and Windows Physical applications, and orchestrates disaster recovery.

These ReliableDR jobs are completely customizable for each unique IT environment, no matter the complexity or mix of virtual and physical machines. This built-in flexibility enables combining physical and virtual application jobs to meet the most stringent compliance requirements. Jobs can be as simple or complex as needed to meet the RPOs, RTOs, and SLAs required for business continuity. For example, you can set up a simple job to certify that a single, critical virtual machine can boot successfully at the DR site within a specified recovery objective. A more complex job can be setup to contain multiple virtual and Windows Physical machines, as shown in Figure 5, where the job boots VMs and instances in the desired order (or in parallel), and then runs any number of lower-level compliance tests to certify hosted applications and services.

Jobs can be executed as tests and failovers. Jobs run as tests generate certified recovery points (CRPs) for the protected services. Jobs run as failovers may either orchestrate disaster recovery using the CRPs, perform a failback after the production environment has been recovered, or perform planned migrations between two different sites. Tests can be scheduled or run on demand. Failovers are run on demand.

Confidence that the applications, databases, networks, and web servers for both Windows and Linux machines are functional is critical for mitigating the risk associated with an outage. ReliableDR verifies granular application functionality, as well as system level availability. Seamlessly integrating with Unitrends Recovery Series physical appliances and Unitrends Backup virtual appliances, ReliableDR tests VMware, Hyper-V, and Windows Physical backups, validating that the backup can be used to recover hosted applications within the specified RPOs and RTOs for each application. When a job runs, ReliableDR interacts with instant recovery (VMWare, Hyper-V and Windows physical) from the most recent backup, running compliance tests, including for boot order (including boot dependencies) to verify that the recovered applications are running properly. If the tests are successful, ReliableDR certifies the backups and guarantees that the protected applications are fully recoverable within specified RPOs and RTOs. In the event of an outage or disaster, ReliableDR will use the certified backups to coordinate and assure recovery.

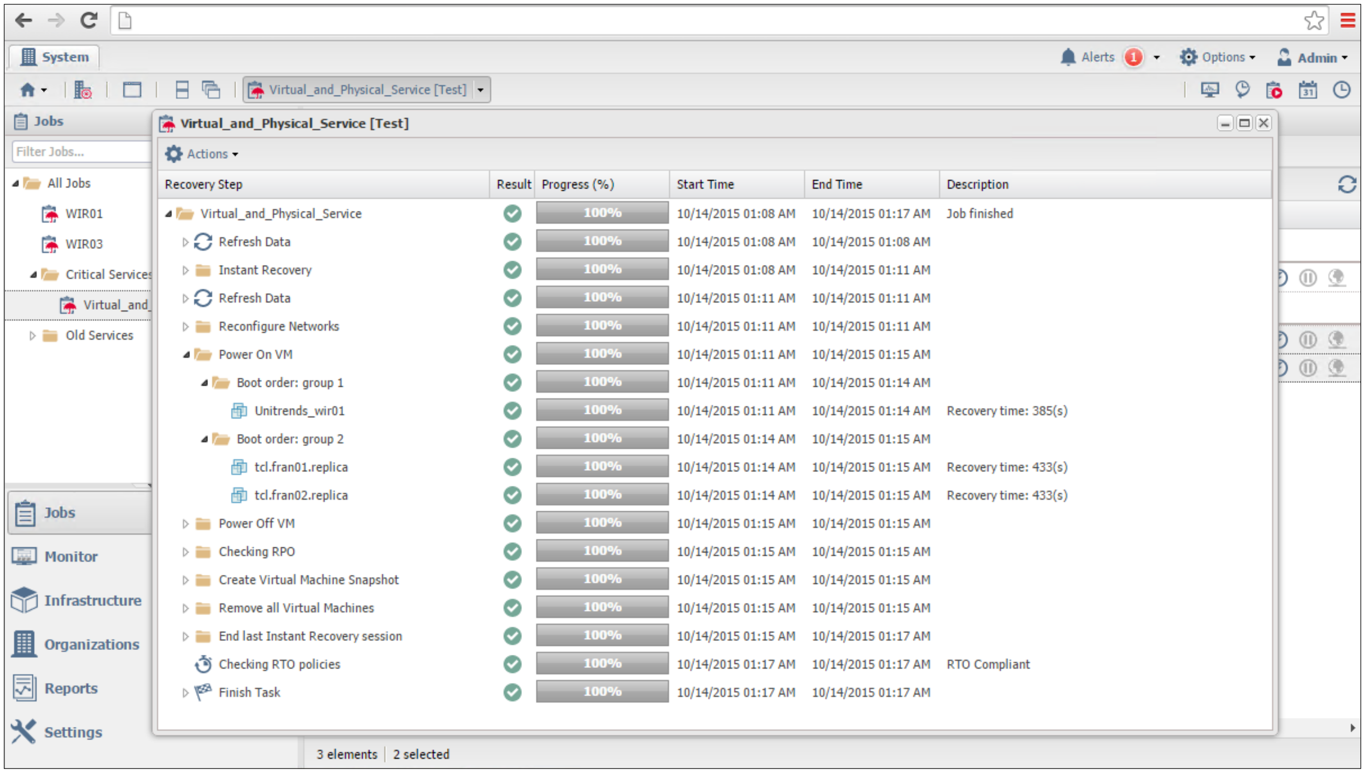

Virtual_and_Physical_Service [Test]

Recovery Step	Result	Progress (%)	Start Time	End Time	Description
Virtual_and_Physical_Service	✓	100%	10/14/2015 01:08 AM	10/14/2015 01:17 AM	Job finished
Refresh Data	✓	100%	10/14/2015 01:08 AM	10/14/2015 01:08 AM	
Instant Recovery	✓	100%	10/14/2015 01:08 AM	10/14/2015 01:11 AM	
Refresh Data	✓	100%	10/14/2015 01:11 AM	10/14/2015 01:11 AM	
Reconfigure Networks	✓	100%	10/14/2015 01:11 AM	10/14/2015 01:11 AM	
Power On VM	✓	100%	10/14/2015 01:11 AM	10/14/2015 01:15 AM	
Boot order: group 1	✓	100%	10/14/2015 01:11 AM	10/14/2015 01:14 AM	
Unitrends_wir01	✓	100%	10/14/2015 01:11 AM	10/14/2015 01:14 AM	Recovery time: 385(s)
Boot order: group 2	✓	100%	10/14/2015 01:14 AM	10/14/2015 01:15 AM	
tcl.fran01.replica	✓	100%	10/14/2015 01:14 AM	10/14/2015 01:15 AM	Recovery time: 433(s)
tcl.fran02.replica	✓	100%	10/14/2015 01:14 AM	10/14/2015 01:15 AM	Recovery time: 433(s)
Power Off VM	✓	100%	10/14/2015 01:15 AM	10/14/2015 01:15 AM	
Checking RPO	✓	100%	10/14/2015 01:15 AM	10/14/2015 01:15 AM	
Create Virtual Machine Snapshot	✓	100%	10/14/2015 01:15 AM	10/14/2015 01:15 AM	
Remove all Virtual Machines	✓	100%	10/14/2015 01:15 AM	10/14/2015 01:15 AM	
End last Instant Recovery session	✓	100%	10/14/2015 01:15 AM	10/14/2015 01:17 AM	
Checking RTO policies	✓	100%	10/14/2015 01:17 AM	10/14/2015 01:17 AM	RTO Compliant
Finish Task	✓	100%	10/14/2015 01:17 AM	10/14/2015 01:17 AM	

3 elements | 2 selected

Figure 5: Mixed Virtual and Physical Job Dashboard

Real-Time Reports

Testing is about risk mitigation. Understanding risk is never complete without being able to see and track recovery results over time for auditors, compliance officers and key stakeholders. ReliableDR provides easy to comprehend visible dashboards and reports that can be sent to key stakeholders or shared through multi-tenant role based access. Showing the test results and how each application complies with its Recovery Time and Point Objectives, displaying both the objective and the measured actual time (RTO Actual, RPO Actual). The RTO/RPO Compliance Dashboard, Figure 6, displays a current view of compliance for each application, green represents successful tests, yellow represents tests that passed with warnings, and orange represents failed tests:

- VMs: The number of VMs included in the job. Clicking on the number will drill into more detailed compliance information for each VM or physical instance. Double-clicking the number will view the job summary.
- Last Test: The days and hours since the last test was run.
- RPO Actual: The actual recovery point and an icon representing the compliance result.
- RPO: The recovery point objective as specified in the service profile selected for the job.
- RTO Actual: The actual recovery time and an icon representing compliance.

- RTO: The recovery time objective as specified in the service profile selected for the job.
- Service Profile: The service profile selected for the job.

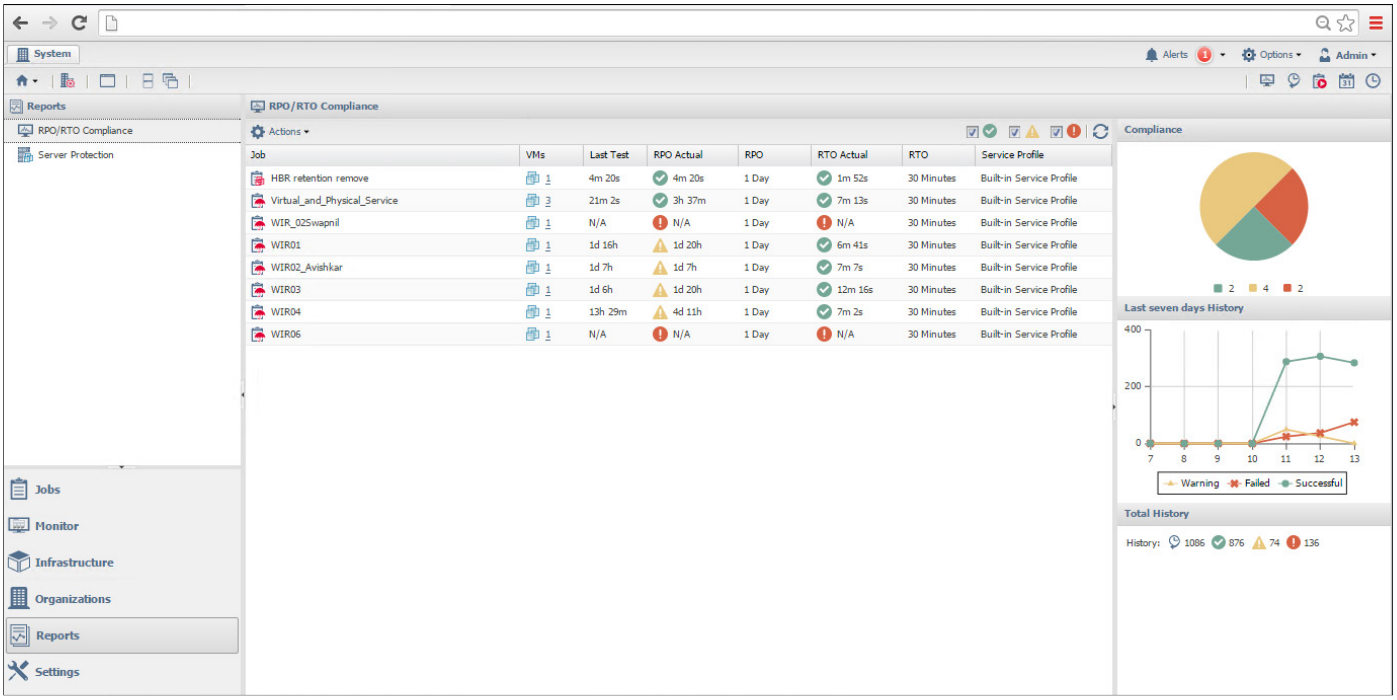

RPO/RTO Compliance

Job	VMs	Last Test	RPO Actual	RPO	RTO Actual	RTO	Service Profile
HBR retention remove	1	4m 20s	✓ 4m 20s	1 Day	✓ 1m 52s	30 Minutes	Built-in Service Profile
Virtual_and_Physical_Service	3	21m 2s	✓ 3h 37m	1 Day	✓ 7m 13s	30 Minutes	Built-in Service Profile
WIR_02Swapnil	1	N/A	! N/A	1 Day	! N/A	30 Minutes	Built-in Service Profile
WIR01	1	1d 16h	⚠ 1d 20h	1 Day	✓ 6m 41s	30 Minutes	Built-in Service Profile
WIR02_Avishkar	1	1d 7h	⚠ 1d 7h	1 Day	✓ 7m 7s	30 Minutes	Built-in Service Profile
WIR03	1	1d 6h	⚠ 1d 20h	1 Day	✓ 12m 16s	30 Minutes	Built-in Service Profile
WIR04	1	13h 29m	⚠ 4d 11h	1 Day	✓ 7m 2s	30 Minutes	Built-in Service Profile
WIR06	1	N/A	! N/A	1 Day	! N/A	30 Minutes	Built-in Service Profile

History: 1086 ✓ 876 ⚠ 74 ! 136

Figure 6: RTO/RPO Compliance Dashboard

A service profile defines the RPO and RTO requirements for each ReliableDR job. To meet the RPO requirement, the applications and all dependencies (example: DNS Services, Active Directory, associated database instances, Web Services, etc.) in the job must be restored to a recovery point within the selected RPO time frame. For example, if the RPO is set to 1 day, the ReliableDR job must restore all applications, including all dependencies to a point in time within the last 24-hours. If any application in the job is restored to a point in time beyond the last 24 hours, the job fails its RPO goal.

To meet the RTO requirement, all applications and their dependencies (example: DNS Services, Active Directory, associated database instances, Web Services, etc.) in the job must be brought back into service within the selected RTO time frame. For example, if the RTO is set to 30 minutes, the ReliableDR job must restore all applications and dependencies within 30 minutes of the job start time. Using service profiles, RPO and RTO requirements can be tailored to each job to suit your business needs and environment.

ReliableDR Delivers Recovery Assurance and Peace of Mind

In summary, ReliableDR provides coordination and automated testing of across heterogeneous (Windows Physical, VMWare and Hyper-V) environments making Disaster Recovery Assurance possible, feature include:

- Guaranteed failover of Windows Physical, VMWare and Hyper-V applications across cluster/clouds.
- 100% automated DR testing; continuous proof that systems are fully recoverable.
- Automatic certification that SLAs (RPOs and RTOs) will be achieved.
- Immediate detection and alerting of deviations in recovery time and recovery point actuals via dashboards, emails and SNMP alerts.
- Application aware with out-of-the-box DR testing for Exchange, SharePoint, SQL Server, MySQL, Oracle, Apache, and any application reachable through http; supports Windows Server and Linux VMs equally.

- Seamless integration with Recovery Series and Unitrends Backup appliance backups and instant recovery for Windows Physical, VMware, and Hyper-V validating recovery of applications local (on premise) or in the Unitrends Cloud integrated with DRaaS for application level disaster recovery assurance.
- Fully automated DR testing dramatically reduces the time, expense and resources required to ensure business continuity while mitigating planned and unplanned outage.
- Real-time and schedule reporting provides audit proof, meets regulatory compliance requirements and automatically submits reports to appropriate stakeholders.
- Scales from small IT to large environments.
- Zero footprint in your production site—sandbox isolated testing eliminates performance and production risk or disruption.

Conclusion

For many years, organizations thought about business continuity in much the same way they thought about business insurance—yes, it was important, but rarely was it top of mind. But that's all changed. Many organizations have, unfortunately, discovered that even a few minutes of service downtime can have lethal effects on their business operations, resulting in lost revenue, diminished customer confidence, and heightened compliance risk.

For those and other reasons, IT executives have raised the bar on business continuity preparedness for their organizations. New technologies, business processes, and partnerships, combined with a raised level of importance for testing and a full appreciation of what virtualization can and can't do for business continuity, are essential to new thinking around avoiding the impact of an unplanned service interruption.

When the real world costs of unplanned downtime are taken into account, recovery assurance is a cost-effective strategy for protecting businesses from downtime, loss of revenue, and reputation.

REQUEST A QUOTE FOR RELIABLEDR RECOVERY ASSURANCE IN THE CLOUD.

About Unitrends
Unitrends delivers award-winning business recovery solutions for any IT environment. The company's portfolio of virtual, physical, and cloud solutions provides adaptive protection for organizations globally. To address the complexities facing today's modern data center, Unitrends delivers end-to-end protection and instant recovery of all virtual and physical assets as well as automated disaster recovery testing built for virtualization. With the industry's lowest total cost of ownership, Unitrends' offerings are backed by a customer support team that consistently achieves a 98 percent satisfaction rating. Unitrends' solutions are also sold through a community of thousands of leading technology partners, service providers, and resellers worldwide.