REDUCING
RTOs TO MINUTES

By Rutrell Yasin

How you can assure customers that their applications and systems are fully recoverable

UNITRENDS

WP-2006-ENG-A | 21 MAY 2014

Hurricanes, floods, fires, earthquakes, blizzards, unstable power grids, terrorist attacks, or even a misconfigured computer system can cause major disruptions to business and government operations.

Whatever the cause of the outage, business and government leaders want to know how quickly operations can get back up and running. With so much of the nation's commerce and government operations running on networked systems, devices, and platforms, these leaders want assurance that their disaster recovery procedures will work when they are really needed.

To achieve this goal, frequent testing of critical applications is imperative. Testing is the only way organizations will have a clear sense of how quickly applications, networks, and systems can be restored after a disruption, otherwise known as the recovery time objective (RTO). The RTO is a service level within which a business process must be restored after a disruption or disaster. Customers must know any business interruption can be kept to a minimum and that recovery processes will work one hundred percent in the event of an outage. Once an organization sets RTOs through a business impact analysis plan, it's up to the service provider to implement an infrastructure that has sufficient resiliency and reflects the recovery time frame that is required by the business. But it's not always an easy task.

Many organizations set RTOs that aren't attainable and do not consider the preparation time, expense, and work disruptions that are typically part of disaster recovery testing. What's more, some organizations only test their disaster recovery (DR) processes once or twice a year and are hampered by complex, manual-driven testing. This means it can take days to obtain RTO information—heaven forbid that a disruption or outage or mishap caused by a mistake during the testing should occur at this time.

What is needed is a high level of automation and orchestration across the various layers of an organization's service delivery stack, from the network to applications. Organizations need the ability to automate replication and failover of critical systems so DR processes can be automated and tested daily at lower cost than complex, manual alternatives.

A fully automated, disaster recovery failover, failback, and testing plan provides continuous proof that systems are fully recoverable. Then, the entire process can be tracked and measured against recovery time and point objectives to ensure that the recovery will be successful and meet service level agreements.

From manual to automated DR testing

Automated recovery requires a host of technologies working in sync, including virtualization, replication, backup, and intelligent storage arrays. Done properly with automated tools, RTO can be reduced to minutes. Plus, service providers can track against objectives with real tests.

Virtualization, which lets organizations run multiple operating systems and applications on a single computer, plays a key role in the move toward automation.

Prior to the widespread adoption of virtualized servers and storage, disaster recovery managers had to copy one-by-one each component running in the production environment and then move them into the disaster recovery site. For instance, if the organization had 10 servers in production, disaster recovery managers or service providers had to map to 10 servers in the DR site. If the organization had six networks in production, well, all six had to be manually mapped to the DR site.

Under these circumstances, recovery was slow and labor-intensive. Storage was backed up periodically. Hardware had to be verified, and operating systems had to be booted individually. Today, most IT infrastructures in businesses and government environments are virtualized to a high degree—around 60 to 70 percent, according to some reports. As a result, service providers do not have to engage in one-to-one mapping any longer.

Virtualization technology incorporates a hypervisor, which lets the tech support team define containers—in which applications can be executed—on the converged machines. These containers can be set up and taken down at will, enlarged or reduced while memory and storage can be increased—all of this can be done without having to buy new hardware.

Additionally, virtualization lays the foundation for cloud computing in which multiple users share the same network services to access computing resources on-demand. Hypervisor containers are essential for multi-tenant, virtualized infrastructures because they let administrators define logical boundaries that separate workloads and data of tenants that share the same hardware. This in turn, leads to a more flexible infrastructure, which in the context of disaster recovery, makes the guarantee of recovery time objectives possible.

Virtual Data Centers On the Fly

How are RTOs more achievable in an automated environment? If a service provider is replicating virtual machines between a production and a DR site, because of the flexibility of the hypervisor, the support team can create on the fly a software-defined data center in the secondary site.

For instance, if a Microsoft Windows server with 8 Gigabytes of memory and 500 Gigabytes of storage needed to be moved, you could request the hypervisor in the secondary site to create one with the same configurations. As the replication is taking place, and at periodic intervals, the replicated bits and bytes can be loaded into the hypervisor container, powered up, and evaluated to determine if all the components come up successfully. The RTO verifies the time it takes to power up the server and the application within it, and this time can be compared with the objective. This gives you a better sense of whether you are complying with your customer's RTO.

Using a DR system that has the intelligence to configure recovery virtual machines in real time, a service provider can replicate the exact configuration of VMs in the production system at the time of the last certified recovery point (CRP). If the primary site's virtual machines change—for example, if memory is increased in one of them—the change is reflected immediately.

Storage hardware and hypervisor components can produce CRPs on the secondary site. Each application can be certified for recovery in a controlled and configured sequence. Snapshots generated by an automated system are consistent and complete, and contain all the necessary components required to restore the entire IT service, including data, operating systems, middleware, and web front ends.

Can RTOs be calculated automatically for hybrid environments since so much of today's critical applications run on a mix of mainframes, Windows servers, Linux/Unix systems as well as virtual machines? Unfortunately, at this time, physical servers cannot be spun up in a software-defined datacenter. However, the good news is that physical servers can be virtualized for DR purposes. Some companies are offering mechanisms that let you backup and then spin up a physical server as a virtual machine in a DR site.

RTOs and RTAs

Another measurement used in conjunction with RTO that you should be aware of is the recovery time actual (RTA). The RTA is established after a disaster recovery test, an actual event, or is based on a recovery methodology developed by the technology support team. This is the timeframe the services provider takes to deliver the recovered infrastructure to the business.

Measurement of the RTA begins at the start of the test and goes to the moment the last component is online and ready to go. RTA is used to measure and compare the recovery objective. As long as the RTA is less than the objective, then everything is fine. If the RTA, for whatever reason becomes longer than the objective, then an alert is sent to the support team notifying them that a component failed to come up or is taking too long.

What is needed is a tool that can calculate how long it will take to recover critical virtual machines and/or their applications. The tool would need a built-in wizard to connect to virtualization technology like VMware. A service provider can then select the virtual machines they want for an RTA estimate, and set the appropriate order for booting them up.

The tool would take a snapshot and create linked clones for each VM. Such a tool can then start up the VMs and time the process, calculating the total time it will take to recover that grouping of VMs. This should give an accurate RTA that can be compared to the RTO and help determine if you can adhere to your customer's SLAs.

4 Steps Toward Meeting Recovery Time Objectives

Having trouble obtaining recovery time objectives that meet your customers' service level agreements? Here are 4 practical steps you can implement to help customers meet their recovery objectives.

1. **Benchmark yourself:** use scripts or free tools in order to understand your approximate recovery time actuals (RTAs) for each of your applications, including Tier2 and Tier3. Review these RTAs with your customers to understand whether they would fall within their expectations.

2. **Understand dependencies:** document the critical servers that support your customers' Tier1 applications and put a process in place to review these dependencies regularly, at least monthly. Misunderstood or new dependencies account for the majority of "gotchas" in DR tests.

3. **Storage and network architecture:** RTOs can be hugely improved by having stand-by replicas in production-grade storage in the DR site. Low-tier applications with longer RTOs can be recovered from backup storage. Ensure that your DR storage, replication processes, and network bandwidth are dimensioned appropriately for the RTO expectation of your application tiers.

4. **Increase DR testing frequency and scope:** resist the pressure to test fewer applications and less often, as this increases your DR risk dramatically. Ensure that new applications always get tested in your next DR exercise, even if they are a lower tier. Leverage automation to increase the testing scope in every exercise, and continuously verify that your RTAs are within business line expectations.

Disaster Recovery Scenario

To give a sense of how this all can be achieved in the real world, take for example, a large insurance company with hundreds of agencies, millions of customers, and thousands of employees. The company needs to ensure consistency and immediate availability of virtual machines and storage arrays in case of a disaster.

The company adopted virtualization technology early, and now has a production data center and secondary disaster recovery data centers. Each has eight x86 farms hosting critical applications; storage is centralized in each datacenter using HP arrays. By deploying the HP array native snapshot and replication capabilities, the technology support team can regularly copy the production virtual machines to the secondary datacenter.

An automated DR tool lets the support team orchestrate VMware hypervisors and the HP storage arrays. To ensure that everything continues to run smoothly in the event of a disaster,

the automated tool 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, the company runs additional recovery point certification jobs every night between midnight and 4am. These jobs create virtual datacenters 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.

DR orchestration takes place out of the secondary datacenter, where the automated DR system runs as a virtual machine. In case of a disaster, at the entire primary datacenter or a subset, IT services can be brought up selectively and automatically by the automated system. This is one illustration of how an automated DR testing tool provides consistency and immediate availability of virtual machines and storage arrays in case of a disaster.

If you are looking for disaster recovery software that automates the replication and failover of your customers' critical systems, Unitrends ReliableDR will let you automate DR failover, failback and testing. The DR software measures accurate RTAs and RTOs of the entire service you are protecting, eliminating time-consuming manual procedures. You can also leverage out-of-the box testing capabilities for each component of your customers' multi-tiered applications to ensure they all function together as a cohesive service at the DR site. An outage or disruption of service is the wrong time to learn that applications cannot be recovered.

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.

7 Technology Circle | Suite 100 | Columbia, SC 29203
866.359.5411 | sales@unitrends.com | www.unitrends.com