Is Your Primary Storage DR Ready?

Storage Switzerland, LLC

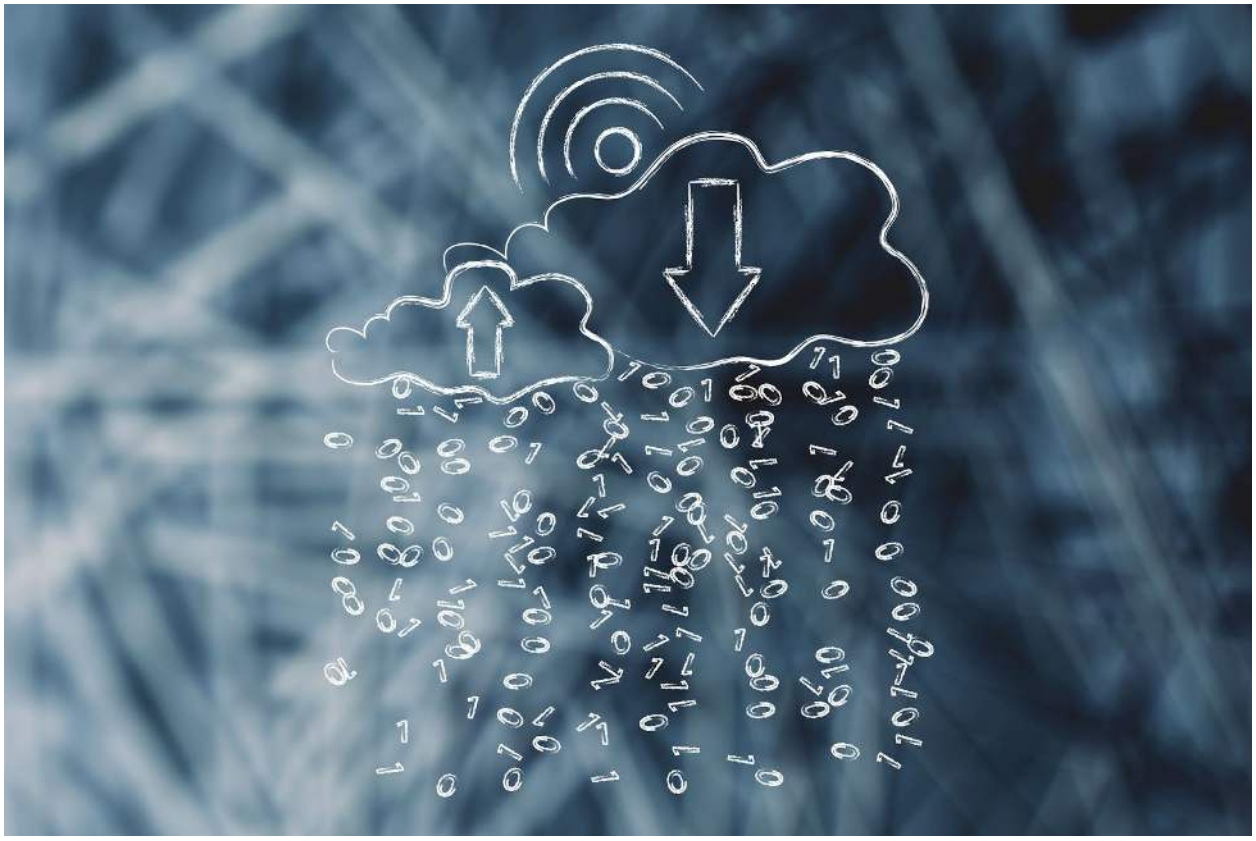

Table of Contents

Chapter 1: Is Your Primary Storage DR Ready?

by George Crump, Lead Analyst

User and organizational expectations are higher than ever. They want high performance access to massive amounts of data. In the event that something goes wrong, they expect IT to resolve the situation instantly and cost effectively no matter how severe the problem might be. To keep pace with these expectations, IT professionals are using a number of data protection schemes ranging from replication to snapshots to more traditional backups. Despite their best attempts, IT is falling behind. IT needs help! The best source of that help is the primary storage system itself, and it needs to be DR Ready. It seems counter intuitive but let's find out why.

What is DR Ready Storage?

Storage that is DR Ready aids in the data protection and disaster recovery (DP and DR) processes without compromising traditional primary storage performance and features. It does this through a combination of its own capabilities (snapshots, basic replication as well as multi-site and synchronous replication) and by complementing or enhancing existing backup solutions – but not replacing them. While many storage systems have replication capabilities, DR Ready storage differs in the implementation of those replication tasks, the number of locations that can be a part of that task and the monitoring of those tasks to ensure successful completion.

Simplicity is Key

A first requirement of DR Ready Storage is to simplify the process – organizations want 100 percent assurance of its recoverability status 100 percent of the time – yet IT professionals are stretched too thin to have time to deal with DR solutions that are difficult to implement. In addition, and potentially more importantly, these solutions need to be simple to monitor in order to make sure the disaster recovery sites updated quickly to meet recovery point and recovery time objectives.

Move Beyond the Basics

The second requirement of DR Ready Storage is it more than covers the basics. Features like snapshots, clones and replication have been available for years. The DR Ready system should not only have these capabilities, it should enhance them. There should be support for a high number and frequent snapshots. Within virtualized environments the primary storage system should enable instant recovery of VM as well as file level recovery within a VM. It should also integrate with the data management capabilities provided by the hypervisor vendors that make local and remote recoveries more seamless.

The replication feature for many primary storage systems is surprisingly limited, often offering only asynchronous replication to a single secondary site. The secondary site is updated periodically based on a snapshot schedule. A few vendors support replication to multiple sites and that should be a required capability. The goal of the multi-site replication is to further insulate against a broader range of disasters or to distribute parts of the workload to multiple offices.

Data centers with extremely strict RPOs and RTOs have an entirely different need, very rapid recovery with zero data loss. Meeting these objectives requires synchronous replication, a capability more rare than multi-site replication.

Another basic to cover is the level of granularity at which the storage system can operate. Many storage systems are limited to the LUN or volume level. Since it is common for virtualized environments to have many virtual machines per LUN/Volume the DR Ready storage system needs to define policy at the VM level not the LUN/Volume level. This VM awareness brings great flexibility to the organization and can speed DR readiness as well as reduce DR site costs.

Integration with Existing Protection Processes

Finally, DR Ready storage has to be careful not to overreach and attempt to replace existing backup solutions for reasons like different storage media, database log management, long-term data retention and more. The dividing line between the two is becoming a bit fuzzy. Generally, primary storage is focused on recovery of the most recent copies of data. However, snapshots and replication can also fulfill the point-in-time capabilities of backup solutions. That said, backup still has a role to play in the recovery of old versions of files and as a recovery of last resort.

It is time for primary storage to do more of the heavy lifting when helping organizations recover data, especially from a disaster. While there are third-party solutions you can add to a primary storage solution, these often increase the cost and complexity of the recovery process. Instead IT professionals need primary storage that integrates advanced disaster recovery techniques making DR preparation, monitoring and testing something that the IT generalist can manage.

Chapter 2: DR Ready Storage Should Complement Backup

by Curtis Preston, Senior Analyst

A popular trend is to have data protection begin where the data begins – in primary storage. But the questions are: Should protecting your data be the responsibility of your primary storage product, or should there be a separation between primary storage and secondary storage and data protection? The answer is both yes *and* no.

Data protection has come a long way in the last 30 years. There was a time when we expected absolutely no data protection capabilities from our primary storage. Thirty years ago, primary storage meant a single spinning disk drive that could fail immediately and take all data with it. The backup system was a completely separate entity. The advent of RAID and redirect-on-write snapshots changed everything. We began to expect that our primary storage would at least assist in helping keep our data safe. RAID protected against device failure and snapshots protected against logical corruption and accidental deletion. Replication of those snapshots to another location provided another level of safety, and we began to have our first DR Ready primary storage.

A primary storage system could – with enough management and reporting system – take over the tasks traditionally performed by the backup and recovery system. Replicating those snapshots off-site can allow the same system to satisfy disaster recovery needs. The ability for a primary storage system to be able to complement the backup and DR systems is good. But it's important to realize that data protection is a spectrum of things that include protections beyond that provided by snapshots and replication. For example, what about long-term storage (i.e. archives) and electronic discovery (i.e. e-discovery)?

The data protection features of primary storage systems should also be integrated into the backup system, so an environment can use these features without abandoning their backup system. Many modern backup systems can schedule the creation and replication of snapshots, as well as catalog the contents of those snapshots. That way you can continue using the backup system that you know and are familiar with, while augmenting its functionality with more modern data protection features from your primary storage system.

Depending on the type of company and its needs, there may be the need to store many years of data. Depending on how that data is to be used, it should be searchable by more than what a traditional backup system or snapshot system is designed to do.

There is archive software and e-discovery software that is specifically designed for this purpose, and until primary storage systems start adding in those types of features, long-term storage needs are probably best served by system specifically designed for those needs.

In addition to being more suited to long-term storage due to searchability and accessibility requirements, there is also the question of cost. Primary storage systems are definitely more expensive than secondary storage systems. When you are considering storing data for many years, cost becomes paramount. While a primary storage system could be used to store data for many years, the cost of doing so on such a system doesn't seem warranted. If a secondary storage system, such as object storage or tape, can meet the requirements of the inactive data access associated with an application, the cost of these systems is more in line with storing data that is not directly contributing to the bottom line.

Many people now take for granted that their primary storage will assist them in protecting the data stored upon it. It's a good thing data protection is now thought of as an integral design component of primary storage. From media protection techniques such as RAID or erasure coding to data protection techniques such as snapshots and replication, data protection is finally getting the attention it deserves. But some functionality, such as long-term storage and archiving, are probably still best left to systems specifically designed for that, both from a functionality and cost perspective.

Chapter 3: DR Ready Storage Should Be A Scalpel Or A Saw?

by George Crump, Lead Analyst

A real disaster is the ultimate test of potentially years worth of planning. Disasters are unnerving because the IT professionals executing the recovery processes have the eyes of the entire organization and its customers on them. IT, which normally operates behind the scenes, is now front in center and the pressure to rapidly recover the organization to an operational state is immense. IT needs focus, not distraction. The problem is most disaster recovery solutions add to the distraction, not lower it.

The modern data center is no longer a one app - one server - one LUN design. Many virtual servers now share a single server and a single LUN or Volume. Most storage systems, though, operate at a volume or LUN level. The lack of granularity creates a multitude of problems, making it difficult to isolate VM for performance guarantees or problem resolution. But the lack of granularity is especially problematic in terms of data protection and disaster recovery.

All VMs are not equally valuable to the organization. Some are mission critical, others are important and some are nice to have. Logically, you apply different levels of protections per importance. Pragmatically, the only applications the organization needs in a ready-state at the DR site are the first wave of apps they must have to return the organization to operation, typically these apps are labeled as mission critical. The problem is that without VM granularity data protection processes like snapshots, backup jobs and replication have to treat all the VMs on the volume the same. It can't differentiate between mission critical, business important, nice to have and not needed at all. The storage system just "sees" a bunch of ones and zeros.

Without VM granularity IT has to apply a single snapshot schedule to those VMs. Then the backup process has to backup all of those VMs. And, yes, while changed block tracking thins the amount of data, all the VMs have to be analyzed for change, wasting precious backup window time. Finally, all the VMs need to be replicated to the remote DR site and, again, a storage system without VM granularity has to replicate everything on the volume to the DR site. This forces IT to filter through the VMs to recover the right VMs.

A second problem with limited VM visibility is understanding VM interrelationships. For example many applications will need a DNS and a Directory server to be restored and running prior to starting. If the data for DNS is on Volume B and the data for the directory is on Volume C but the organization is only replicating Volume A (where the application is) then the DR attempt may fail or at least be slowed while DNS and Directory servers are restored from backup.

A third problem is cost. For most organizations a full scale disaster will be a rare experience. Obviously, the impact of that rare experience occurring is grave enough that it does have to be planned for, but its rarity also means that being wise instead of speculative or naïve with DR site expenditures also makes sense. This is especially true when considering that the first wave of recovered applications is relatively small, which makes it even more expensive. To make the problem worse, often the cases being dealt with are human errors batching deleting some critical VMs. The alterative to handle these “not-so-rare” scenarios in a cost-effective way becomes increasingly important and visible within an organization.

Empowered by a VM aware storage system, organizations can design the DR site to be able to handle only that first wave. All subsequent recoveries, for business important or nice to have applications, are recovered as additional equipment is ordered in. The value of the storage system understanding the VMs that it stores is that only these mission critical VMs need to be replicated and be stored in a ready-state at the DR site. The rest can be protected by the backup application on cost effective backup storage or even tape. As a result VM granularity leads to the purchase of a smaller and faster responding storage system and less server hardware at the DR site.

Successful recovery from a disaster requires a scalpel not a saw. Being able to surgically identify just the VMs needed to return an organization to operation is vital to focusing IT responders on the task at hand. At the same time this precision keeps costs of the DR site under control. Most mission critical workloads are now virtualized and intermingled with less important applications and data. Storage systems with VM understanding allow for easier identification and isolation of the most important VMs during a disaster and are critical to successful and affordable DR.

Chapter 4: DR Ready Storage Must Be Space Efficient

by W. Curtis Preston, Senior Analyst

One of the biggest criticisms of using primary storage for DR is cost. Primary storage can be expensive, and using it in your DR plan essentially doubles your storage requirement. Therefore, it is essential that DR Ready storage use as little space as possible and provide granular protection policies. Let's take a look at the ways that it can do that.

One of the primary ways to reduce actual storage use is to use thin provisioning. This feature allows customers to provision much larger volumes than they actually need, while the storage product only uses blocks that they actively use. For example, a user can create a 10 TB “logical” volume that only contains 1 TB of “actual” data. This feature saves space, aka cost and makes management easier, as storage administrators do not have to continually re-provision their volumes as they need more space.

The next feature that is important, both for data protection and for space efficient use of storage, is snapshots. Snapshots are an essential part of storage and often combined with replication for disaster recovery. Replicated storage without snapshots can be just as useless in a disaster as no storage, because the replication would also replicate any kind of logical corruption such as a black hat attempting to damage your data. Snapshots allow you to go back in time prior to any malware or other corruption to your system.

A snapshot – as the name implies – is a *view* of your storage rather than an actual copy. Subsequent snapshots take up very little space, as they only save the changed blocks and pointers since the original snapshot was taken. This is very good from a storage efficiency perspective, as one can have hundreds of snapshots without requiring a significant amount of additional storage to store them. In fact, storing 100 snapshots takes no more storage than storing a few data changes, because each snapshot only stores the blocks that are unique to that snapshot.

The one thing that one must keep in mind about snapshots is each snapshot is a virtual copy. Snapshots must be replicated to another destination in order to be able to use them in a disaster. A snapshot of your storage after it catches fire is as useful as a snapshot of your house after it catches fire. You must replicate snapshots for them to be useful in DR.

As we mentioned earlier, the DR Ready system should provide administrators with granular control over which of these snapshots gets replicated because of the additional efficiency and cost savings gained when done at the right abstraction level.

Two other technologies are compression and deduplication. They are mainstream now but confusion and hype still remains. It's important to realize that the absolute space savings factor is not the full picture. Their savings differ when applied to different workloads.

Compression looks for redundant patterns within a given file or volume, and replaces those patterns with pointers. The easiest way to understand compression is to consider a large document that happens to have the same word repeated many times. A compression algorithm would notice the repeated word and replace all of the occurrences of that word with a single reference, saving a significant amount of space. If, however, that same word was repeated in *another* word processing document, compression would not notice that.

Deduplication, on the other hand, looks at patterns *between* volumes or files. In the example above, deduplication would notice that the same word is present in multiple files, and would replace all incidents of that word with pointers. The most applicable use of deduplication in primary storage is when storing VM images. Each VM image contains a copy of the operating system and one to many applications. These redundant parts of the VM can be identified and replaced with pointers, resulting in significant space savings for those blocks.

Each of these technologies have their role, and save a significant amount of space; however, mileage will vary depending on the workloads. DR Ready storage needs to counteract the cost of having a second storage system that is used only during a disaster. When all of these data reduction technologies are implemented, it is possible to have DR Ready storage with similar cost structure as maintaining an entire backup and recovery system. And there is no question that DR Ready storage is a much nicer thing to have than a traditional backup system.

Chapter 5: DR Ready Storage Must Be Cloud Aware

by George Crump, Lead Analyst

The success of cloud providers like Amazon, Azure and Google is forcing the data center to re-think various aspects of their IT infrastructure. There are cases where a cloud provider can enhance typical on-premise storage to be DR Ready and make long term retention viable.

Most organizations hope to never or at least rarely need archives. However, because of the implications to the business of not preserving their data, it is something they have to invest in. That investment may include development of another facility, along with the effort of purchasing or renting the physical location and then putting in all of the equipment.

The cloud has the potential to lower the "rent" and eliminate acquisition costs. However, there are cautions.

Data Has Gravity

It might be easy and even free to import data into public cloud but could cost a great deal to try to get any data back. On top of that, you are compromising on the level of control by choosing to pay the “rent.” IT professionals need to weigh the costs of quick startup vs. the long term costs of renting storage. IT also has to be careful before jumping on the cloud DR bandwagon, the cloud’s DR site capabilities are not yet up to 99.9999 percent SLA required by many enterprises. The key is flexibility, a DR Ready storage system should be cloud-aware so the organization can choose to use the cloud in a way that makes the most sense for it.

While many on-premises storage systems claim to have the ability to integrate with a cloud provider, it is a kludgy afterthought, forcing IT professionals wanting to leverage cloud economics to find a third party solution, even though their existing storage system is oh-so-close to being cloud ready. It’s wise to take a closer look how that connectivity can lead to the savings you originally intended.

What Can You Do With Your Cloud Connection

Every primary storage system worth its salt has snapshot technology and most of those leverage that snapshot technology to replicate data to a disaster recovery site. The problem is the target system almost always has to be a near identical system. It is, of course, unlikely cloud providers will implement one of every storage system into their infrastructure.

What if the primary storage solution has the ability to replicate snapshots to an simple storage service (S3) target? S3 is the API Amazon provides to interface with its object storage, and it has become the defacto standard for object storage interfaces. Almost every cloud provider and most object storage systems are either native S3 or provide an S3 interface.

With an S3 interface some or all of an organization's data could replicate to a cloud provider instead of, or in, addition to a secondary system in another facility.

The Value of Cloud Aware

Once this data is in the cloud, an organization could likely “check-off” its long term retention requirement. However, from a disaster recovery perspective, IT professionals should be more careful. Replicating to a secondary site with a similar system as opposed to the cloud may still be the better option. A secondary site is under the organization’s control and has a known enterprise class storage system. A quality system at one or multiple alternate sites is critical as these systems may one day become the primary in the event of a disaster. In this situation the cloud is still attractive as a location to get data “when all else fails.”

For organizations where a second site is not available, the cloud as a DR site remains an option. Assuming IT works out the obvious network configuration changes, recovery from a disaster could occur quickly without having to wait for the implementation of new equipment or storage. Organizations will have to transform their VMs from the hypervisor they use on-premises to the hypervisor the cloud provider uses. While the organization saves the upfront costs of having to physically buy facilities and physical equipment, there is more work required to ensure that the failover process will work and is understood.

While cloud integration is important, “lift and shift” your entire DR to cloud is another extreme. Organizations with secondary sites available to them may want to look at those sites as the first option and the cloud as the recovery point of last resort.

Long term data retention has become one of the best use cases of the cloud. The economic and logistical advantages are very compelling. Using the cloud for DR is another option especially for single site organizations. Organizations with multiple sites should look for storage systems that provide the ability to replicate to those sites while also being cloud aware. Then these organizations can use the cloud for long term data retention and as a backup to the DR site. In the end it is all about flexibility. Your primary storage should provide the flexibility to leverage and integrate with the cloud for both DR and long term retention models.

Chapter 6: DR Ready Storage Must Perform

by George Crump, Lead Analyst

The ability to quickly resume operations after a severe disaster is critical for organizations of all sizes. Fortunately, disasters don't occur every day, in fact they are actually pretty rare.

It is their severity that makes organizations plan for them. The request to improve application response time on the other hand, is a daily demand, as is the constant pressure to reduce costs.

The recent Amazon Web Services outage also taught us that disasters are often caused by innocent human errors that the infrastructure needs to be prepared for. DR Ready Storage needs to meet the day to day demands as well as help the organization withstand the worst case disaster.

The Problem with the Just Throw Flash at it Mentality

The first response to a request for more performance is to leverage flash media. But flash storage, without quality surrounding technology, won't perform to its full potential nor will it be used efficiently. The software that drives the storage system should be flash aware so its performance can be fully exploited. It should also provide detailed analytics so performance sensitive applications are given priority. The problem with most storage analytics, if the storage system even provides them, is they are not granular below a volume or LUN level. This means in a virtualized environment it is impossible to differentiate and prioritize specific VMs. Deploying flash for the sake of just a speed increase by itself is like upgrading a race car to Formula 1 hardware standards but keeping the driver maneuvers at amateur level.

Managing Scale

Another storage infrastructure challenge is managing growth. As the environment continues to scale most environments will need to expand their storage systems to meet either a need for more capacity, performance or both. Scale-up architectures require that the system's controllers receive an upgrade, or add another unit, increasing management expense. Scale-out system address the scale-up shortcomings by aggregating capacity and performance from across multiple nodes but these systems often require three nodes to get started, tightly coupled federation among the nodes via a complicated clustered file system, plus they often don't fully utilize the capabilities of each node, lowering efficiency.

IT professionals should look for storage systems that "scale-right."

Meaning they can purchase the product as a scale-up system – one node to solve a specific problem – then gradually add nodes as more performance or capacity is needed but have those nodes loosely coupled, scale independently from compute, and all managed from a single management interface. Once again VM granularity comes into play here, facilitating the ability to move VMs between storage systems, automatically and non-disruptively based on policy.

A Role for Hybrid?

While the primary data center is becoming increasingly all-flash, there is still a role for hybrid storage, and there are advantages in working with a vendor that can provide both. A hybrid system is an ideal target for the DR Ready storage system, especially if the target's sole role is disaster recovery. In this case it may make sense to use a hybrid system, with a small flash tier and a large hard disk-based tier. If there is a failure, mission critical VMs, which will be recovered first, will automatically be promoted to the flash tier ensuring excellent performance, even in the recovered state. One thing to note, the key factor here is nothing nostalgic about spinning disks vs. solid-state arrays. It all boils down to the balance between performance and cost, which take on different profiles and when the crossing point might occur.

Recovery from a disaster is, of course, critical and DR Ready storage makes that so much easier to achieve. But at the same time, organizations have to deal with the day to day demands of the business and near the top of the list is the ability to provide excellent application performance and to be able select a storage system that can grow as the organization grow. However, the art to balance the performance requirements and disaster recovery needs should not be limited to just compromises. DR Ready storage should take into account the moving cross points of performance and cost profiles, and accept the change as the constant norm.

Chapter 7: Successful Disaster Recovery? - Trust But Verify

By W. Curtis Preston, Senior Analyst

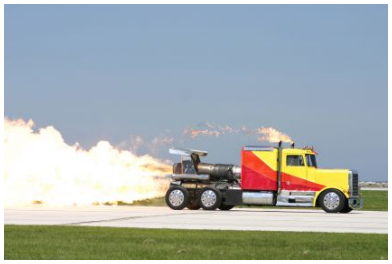

Verifying that you're ready for a disaster is difficult but not impossible. The three elements of your infrastructure that must be present for a successful recovery from a disaster are compute, network, and storage. Let's take a look at how you prepare each of these elements for disaster.

Making sure you have adequate computing resources in a disaster is significantly easier today than it has in days past, thanks to the invention of server virtualization and cloud computing. In the old days, one had to provide their own servers during a disaster, or contract with another company to provide the servers. The difficulty of matching production server hardware with DR server hardware was a constantly losing battle, as DR hardware had to be updated every time production hardware was updated. Now that both sides of the equation are using server virtualization, updating the DR configuration costs nothing but the time spent using the configuration tool to update the VM configurations. There are also tools that can completely automate that process.

Similarly, making sure you have the appropriate network infrastructure for disaster is a lot easier than it used to be. Most infrastructure providers have much more bandwidth available than a typical DR customer would require in order to perform their job. Therefore, the hardest part of making the network infrastructure ready for disaster is setting up the DNS and VPN configurations so that the DR network can behave as if it is part of the local data center. This is not a simple process, but the IT administrator can automate it.

Preparing the storage infrastructure and the data that resides on it is also significantly easier than it used to be with the advent of DR Ready storage. It is a much simpler process to define and make ready storage of equivalent capacity and performance. It's also very simple to continuously update the data on that storage. With modern storage at the right level of abstraction, automation is significantly easier.

Trust But Verify

Just because it is possible to define and make ready all of the elements of the infrastructure in advance, that doesn't mean it will actually be ready in the case of a disaster. The only way to verify that you're ready for disaster is to test a recovery on a regular basis. Unfortunately, most people test only a small portion of their infrastructure when doing a DR test – if they test anything. They recover a single application or database. In fact, many people doing a recovery test perform only a data restore or verification. They do not actually restore an entire application.

It's important to understand that modern applications usually use a variety of resources residing on several interdependent systems. Where historically you could restore a single database and know the application using the database was on the same system, this is no longer the case. So proper DR testing must first acknowledge many systems are related and they must be recovered in groups. These groups are referred to as *consistency groups* because the data between different VMs must be from a single point in time or it is not consistent. You cannot recover two different databases to two different points in time and then expect them to work together without some type of referential integrity issue.

This is why DR Ready storage needs to be able to understand the data being stored, and understand the concept of consistency groups so that all related data can be restored to the same point in time. Any DR tests should be restoring consistency groups, not simply a single database, application, or file system. IT personnel must be aware of this and trained on this by performing frequent DR tests. Being ready for disaster is as much about preparing your personnel and automation process as it is about preparing your infrastructure. If the first time your personnel are using your DR infrastructure is during a disaster, it's going to be a disaster.

There is no reason a modern company should suffer a major outage during a disaster. But the reality is far from ideal. Stories abound of companies that do not survive major outages. Don’t let that be your company. Contract cloud services that can be made available in an instant in case of DR and test them upfront. The technology is there and the cost is reasonable. Just make sure to verify your trust in these service, especially since automation and data awareness make verification much easier and less expensive than it used to be.

Chapter 8: Tintri Enterprise Cloud Platform is DR Ready

By Joseph, Senior Analyst

Chapter 1 discussed the need to simplify data protection and disaster recovery operations, with primary storage that is DR Ready. DR Ready primary storage would aid in DP and DR processes without compromising traditional primary storage performance. A good example of a comprehensive DR Ready storage system is Tintri's all-flash and hybrid arrays.

The Tintri solutions are a scale-out storage platform for virtualized and cloud workloads that meets the expectations of organizations needing greater levels of availability and recoverability. Tintri offers a primary storage system with fully integrated DP and DR. It also integrates with existing DP and DR solutions like Commvault, Veeam and to the public cloud using S3 API based connector. The connector allows Tintri snapshots to be stored to clouds like AWS or to on-premises S3 compatible object storage.

Tintri Hardware

Tintri offer two types of storage arrays. The T5000 series is an all-flash array, while the less expensive T800 series is a hybrid array.

The T5000 all-flash series can scale up to 308TB for 5000 VMs in a single 2 RU chassis and uses high density 3D NAND drives and multiple capacities with expansion options to provide flexibility, performance and capacity. This can scale out to 10 PB of all-flash storage in 1.5 racks, supporting up to 160,000 VMs. Tintri's storage delivers performance isolation for every application and policy based per-VM quality of service (QoS) while ensuring low latency and high performance across the entire flash capacity.

The T800 hybrid series can scale up to 120TB for 3500 VMs in a 4 RU chassis.

All Tintri storage arrays use Tintri OS, which is a virtualization-aware file system. It offers VM-level QoS, analytics, data management, and automation. It supports vSphere, Hyper-V, RHEV, OpenStack, and XenServer hypervisors.

Management and Analytics

All Tintri operations are done through the Tintri Global Center (TGC), which allows access to all features, functions and configuration options. TCG gives access to powerful analytics capabilities that let you see all aspects of replication jobs (both current and historical), across infrastructure including host, network and storage (both local and remote). The advanced version also allows you to pool the capacity of multiple Tintri storage nodes and can be balanced based on workload requirements.

DP and DR Capabilities

Tintri storage software provides robust DP and DR features to help ensure the safety of the stored data while also facilitating rapid data recovery in the event of an outage or system failure. Tintri executes all of these features at the individual VM level. The DP and DR features are:

- Snapshots and clones at the VM level. This allows recovery of the full VM, a single virtual disk, or folder or file, with a single click. You can also clone a replica into a live VM. This can facilitate and simplify DR testing.
- Asynchronous replication: one-to-one, many-to-one, one-to-many, allowing up to four different destinations. Policies are applied at the VM or service group level, sending only changed data that is deduplicated and compressed.
- Synchronous replication: Zero Recovery Point Objective (RPO) and near zero Recovery Time Objective (RTO) ~30 seconds with one-click failover at the submount level.
- Snap recovery and File Level Recovery (FLR)
- VMware SRM integration
- RESTful API, PowerShell and vRO plug-in simplify and automate workflows

The Tintri OS provides great flexibility allowing asynchronous, synchronous and non-replicating VMs on the same array as well as supporting bidirectional replication. Replication can take place between different types of Tintri arrays, both all-flash and hybrid. Many other solutions require the exact same type of array for both source and target. Tintri does not have this limitation. The ability to use less expensive target arrays can significantly reduce costs at the DR site.

Tintri has a powerful primary storage solution that meets the DR Ready standard. It provides functions and controls at the individual VM level rather than at the LUN level. This allows for selective replication of VMs that require it rather than all the VMs on a given LUN. Additionally, in a recovery situation, you can recover selected VMs without the need to wait for the recovery of the entire LUN or volume. It also has excellent DR features that greatly simplify management and recovery operations so you don’t need a DP and DR specialist to get it setup and manage it.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the storage, virtualization and cloud marketplaces. Our goal is to educate IT Professionals on the various technologies and techniques available to help their applications scale further, perform better and be better protected. The results of this research can be found in the articles, videos, webinars, product analysis and case studies on our website **storageswiss.com**

George Crump, Chief Steward

George Crump is President and Founder of Storage Switzerland. With over 25 years of experience designing storage solutions for data centers across the US, he has seen the birth of such technologies as RAID, NAS and SAN. Prior to founding Storage Switzerland he was CTO at one the nation's largest storage integrators where he was in charge of technology testing, integration and product selection.

Curtis Preston, Lead Analyst

W. Curtis Preston (aka Mr. Backup) is an expert in backup & recovery systems; a space he has been working in since 1993. He has written three books on the subject, Backup & Recovery, Using SANs and NAS, and Unix Backup & Recovery. Mr. Preston is a writer and has spoken at hundreds of seminars and conferences around the world. Preston's mission is to arm today's IT managers with truly unbiased information about today's storage industry and its products.

Joseph Ortiz, Lead Analyst

Joseph is an Analyst with Storage Switzerland and an IT veteran with over 35 years of experience in the high tech industries. He has held senior technical positions with several OEMs and VARs; providing technical pre and post sales support as well as designing, implementing and supporting backup, recovery and data protection / encryption solutions along with providing Disaster Recovery planning and testing and data loss risk assessment in distributed computing environments on Unix and Windows platforms.

About Tintri

Tintri provides enterprises and cloud service providers with an enterprise cloud platform that offers public cloud capabilities inside their own data centers. Combining cloud management software, web services and a range of all-flash storage systems, Tintri not only delivers many of the benefits of public cloud infrastructure including agility and automation, but also gives organizations the control and better economics they need to build agile development environments for cloud native applications and to run mission critical enterprise applications. Tintri enterprise cloud helps organizations to improve speed to market and provides IT as a service to internal business groups. That's why leading cloud service providers and enterprises, including Comcast, Chevron, NASA, Toyota, United Healthcare and 20% of the Fortune 100, trust Tintri with enterprise cloud. Visit https://www.tintri.com to learn more.