Disaster Recovery in a Hyperconverged World

By Howard Marks

MARCH 2017 | US$25 | S070317

Channel Partners™

Disaster Recovery in a Hyperconverged World

By Howard Marks

MARCH 2017 | US$25 | S070317

Channel Partners™

Table of Contents

Channel Partners™

About the Author

linkedin.com/in/homarks
@deepstoragenet

HOWARD MARKS is the founder and chief scientist at DeepStorage LLC, an independent test lab and analyst firm specializing in data storage, virtualization and data center networking. Before founding DeepStorage, Marks spent more than 30 years as a New York-based consultant helping organizations including BBDO, SUNY Purchase and the Foxwoods Resort Casino solve their IT infrastructure problems.

An entertaining and highly-rated speaker, Marks appears regularly at industry events including VMworld, Interop, SNW and Microsoft's TechEd. He has written three books and hundreds of articles on networking and storage technologies since testing more than 100 modems for a special issue of PC Magazine in 1987. He also co-hosts the "Greybeards on Storage" podcast.

Disaster Recovery in a Hyperconverged World

DESIGNING AND OPERATING IT INFRASTRUCTURE HAS, AT LEAST SINCE THE DECLINE OF THE MAINFRAME, been a bit of a black art, requiring shamans gifted in the independent disciplines of compute, networking and, most arcane of all, storage. Hyperconverged infrastructure (HCI) changes that by integrating the storage layer into the servers that host virtual machines and management tools that control them. This vastly simplifies both design and operations.

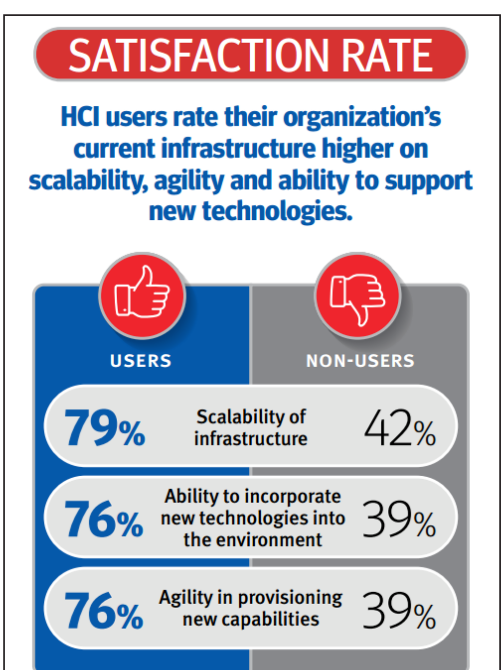

Source: MarketConnections

For partners, hyperconvergence lessens complexity and thus saves money. Gartner says hyperconverged gear will be 24 percent of the integrated systems market by 2019, reaching almost $5 billion in sales, and will be fully mainstream by 2021.

If, as proponents suggest, the HCI model of building a system out of standard building blocks, each of which provides integrated compute and storage resources, is revolutionizing the data center, it is also changing how we must protect customer data and applications from disaster.

Any disaster recovery solution has two basic functions. First, it must replicate data outside the blast radius of any disaster from which your customers might need to recover. Then it has to simplify the process of turning stored data at the disaster recovery (DR) site back into running applications.

Most traditional DR solutions rely on the storage system to replicate data from the primary data center to a DR site. As we shift from monolithic disk arrays to

HCI, we have to remember that the old dinosaur-like array not only provided storage capacity, it also replicated the data it held to its doppelganger at our DR site. How will you replicate that function? As we displace dedicated storage arrays with hyperconverged systems, we must replace not just the capacity and performance the array previously provided but also its replication capabilities.

Some of the more sophisticated HCI solutions, like HPE's SimpliVity and Nutanix, integrate the replication process with their storage software. While the architecture resembles storage array replication, the hyperconverged solutions have a couple of significant advantages.

Storage arrays present for replication logical volumes known as LUNs, for Logical Unit Number, basically the SCSI address of the logical disk a volume represents. Most arrays can handle only a limited number of LUNs. For this reason, and because provisioning and connecting these LUNs to the hosts requires more than a little work, one LUN usually contains 10 or more virtual machines.

Hyperconverged systems, in contrast, manage their storage at the level of the individual virtual machine (VM) — the same granularity they use for compute.

Managing replication, snapshots and other data protection features at the VM level makes it much easier to ensure that each application gets the level of protection it requires. Customers using LUN-level replication must either carefully group VMs with similar protection requirements into common LUNs or replicate all the VMs in each LUN at the replication frequency and retention schedule of the most critical VM stored on that LUN.

Overprotecting less critical VMs requires more network bandwidth between the primary and DR site and more storage capacity for retention.

The VM-centric management model enabled by HCI also simplifies the recovery process. An administrator using storage-based replication would have to mount the affected LUNs at the DR site and then bring the VM images they want to restore in to the management service, like vCenter for vSphere or SCVMM for Hyper-V.

On an HCI system, they could view protected VMs in the HCI management console and restart them with much less effort.

Hyperconvergence and Branch Office Data Protection

HCI may be attractive to customers with remote and branch offices that need local infrastructure. In an ESG survey of **347** IT professionals responsible for ROBO environments, **37%** said data backup is a problem, and **30%** have disaster recovery difficulties. Combined with DRaaS, the simplicity of HCI's "building block" architecture can ease backup and DR at remote sites.

The Replication Hierarchy

The most basic DR solutions are based on the oldest of IT data protection technologies, the lowly backup. While using backup for disaster recovery brings up unhappy memories of sleepless nights restoring servers from tapes in my youth, today's solutions are a lot better. Let's look at a few in the context of HCI.

Backup-based solutions: A Jurassic-era "restore from tape" DR plan can at best meet an RPO of one day; restoring last night's backup means all your customer's servers revert to last night's data, after all. For an RTO, plan on the several hours it will take to restore a 300 TB SQL Server.

Modern backup software, like Veeam Backup and Replication, will back up protected VMs then replicate the backup file to a repository at the DR site. When a disaster is declared, the VMs can be mounted directly from the backup repository.

While these solutions can run the backup-and-replicate cycle as frequently as once an hour, the backup process can impact the performance of production systems, so they're generally used to protect data that can handle reverting to the previous day's state. This coarse granularity combined with the backup application's data reduction capabilities minimizes the amount of data transmitted, making backup-based solutions the most network efficient of all the replication options. That's true on conventional and hyperconverged infrastructure.

Disaster Recovery Objectives

Regardless of the underlying technology, any disaster recovery plan must be designed with three fundamental objectives in mind:

- **Recovery Time Objective.** RTO is the time it takes to recover an application once the recovery process begins. While we casually refer to an application's RTO as the maximum amount of downtime the organization can endure, planners should make sure to factor in any time between a failure and the beginning of the recovery process.
- **Recovery Point Objective (RPO).** RPO defines the time period for which data will be lost if the application needs to be recovered.
- **Distance.** A customer's DR site should be far enough from their primary data center to ensure that the primary data center and DR site are not both taken off line in the same disaster.

Note that these definitions describe the recovery of an application, not an entire data center. Different applications will have different recovery point and time objectives to meet business needs.

Options like Veeam Cloud Connect send replication traffic to solution providers' data centers, allowing those solution providers and their resellers to offer disaster recovery-as-a-service (DRaaS) and freeing user organizations from the overhead of a dedicated DR site.

Point-in-time or snapshot replication: Point-in-time solutions leverage the HCI system's snapshot capabilities by periodically taking a copy and replicating the frozen changes to the DR site or sites. These systems typically allow snapshots to be created as frequently as every 15 minutes; customers can specify how many snapshots to retain at each site. The biggest problem with point-in-time replication is that most vendors use sloppy language and call it "asynchronous replication," causing customers to confuse it with actual continuous asynchronous replication, described below. This sloppiness also leads vendors to claim that a system taking snapshots every 15 minutes provides an RPO of 15 minutes; that ignores the time it takes to replicate the data between sites.

As a rule of thumb, assume that a point-in-time replication system can deliver an RPO of twice the snapshot period. Exactly how long it takes, after the snapshot is made, for the data to arrive at the DR site is of course a function of the network bandwidth between the sites and how much data there is in each snapshot.

Point-in-time is the most commonly supported replication technology in HCI solutions including Scale Computing, Nutanix, Simplivity and VMware's vSphere Replication, which is storage-agnostic. The more sophisticated systems support snapshot consolidation to provide retention schedules that reduce the number of snapshots as they age. They can, for example, store hourly snapshots for one day then prune the data so there's only one snapshot per day for the previous seven days.

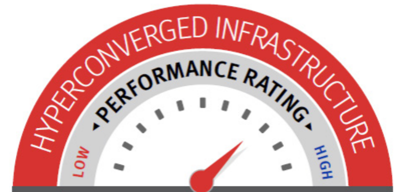

Hyperconverged infrastructure users are more satisfied with key infrastructure functions

NON-USERS	KEY FUNCTIONS	USERS
27%	Data storage, backup and recovery	79%
20%	Application performance	72%
29%	Application uptime	69%
31%	Compliance support	83%
20%	Speed of implementing new applications	76%
24%	Ease of deployment	79%

Source: MarketConnections

Asynchronous replication: Asynchronous replication solutions capture each write request to the storage target and transfer them in real time, and in order, to the DR site, where they are applied to a virtual disk. When the primary data center goes offline, only the data "in flight" from the primary data center to the DR site will be lost, allowing asynchronous solutions to achieve an RPO of just a few seconds, given sufficient network bandwidth.

Asynchronous replication systems acknowledge each write to the application before transmitting the data to the remote site and will journal data for transmission later if the network becomes congested or fails. This journal is usually rather small, holding minutes, not days, of data.

More sophisticated solutions like Zerto's Virtual Replication keep a significantly larger journal at the secondary site and allow users to restore to multiple points in time from the journal. Since application servers frequently corrupt their data as they breathe their last, this journal allows admins to restore to the last good data.

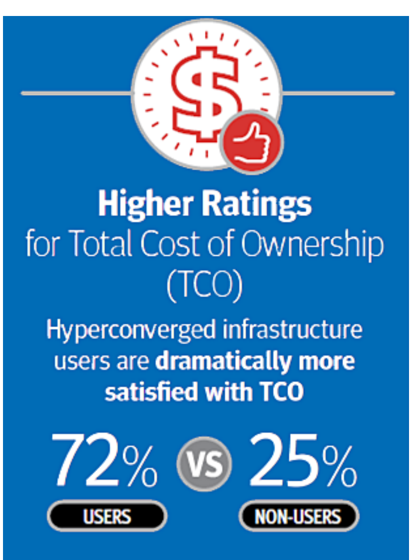

Source: MarketConnections

Synchronous replication: Like the continuous asynchronous replication described above, synchronous replication systems capture and replicate each write request. The difference is that the synchronous replication system doesn't acknowledge writes to the application until the data has been written to both sites.

Synchronous replication therefore achieves an RPO of zero, with no data loss, but this high level of protection comes at a price. Since the application doesn't get an acknowledgement until data is written at both sites, the network's latency, a minimum of 1 millisecond per 100 kilometers of link, adds to storage latency and therefore reduces storage performance.

We're not aware of any HCI solutions providing synchronous replication.

Recovery Isn't Just a Storage Problem

A full recovery involves more than just mounting VM images into a hypervisor and powering them up. The IP addresses, both internal and internet-facing, at the customer's DR site are probably different from those at the primary site, and you'll likely have to readdress all your VMs and update your DNS servers to accommodate the change.

While none of today's HCI solutions automate this process today, several, including Nutanix and Simplivity, support VMware's SRM recovery automation tool.

HCI in the DR Site

Using an HCI solution for recovery makes sense even for customers that continue to run more conventional architectures at their primary sites. The simple building-block expansion model HCI provides makes it easy to add remote "bricks" as the number and size of workloads in the primary data center grows. HCI systems provide for remote maintenance, eliminating trips to the DR site to update firmware or perform other maintenance.

The compute resources provided by HCI solutions make it easy to take advantage of native replication for Active Directory or Exchange DAGs and can provide a prepopulated test/dev environment with the addition of a few VDI instances for developers.

While HCI hasn't yet revolutionized the DR process the way it's changing data center design, HCI is attractive for may DR applications. It will be interesting to see how much better DR support from Nutanix, HPE and Scale Computing will be with another few years' seasoning.

Related Reports

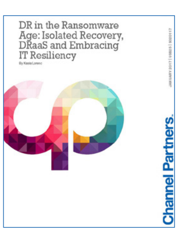

DR in the Ransomware Age: Isolated Recovery, DRaaS and Embracing IT Resiliency

Ransomware attacks are on the rise and are incorporating more complex and sophisticated techniques. Security experts agree: It's not *if* an attack will come, but *when*. Channel partners can help their customers arm themselves against the devastating effects of ransomware with the right tools, infrastructure, approach and strategy. And a key element in their arsenal is one not-so-secret weapon: disaster recovery.

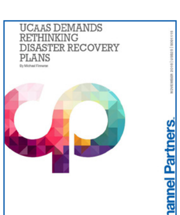

UC Demands Rethinking Disaster Recovery Plans

Now that the migration of voice and UC systems to the cloud has ushered in ubiquitous connectivity, your clients — and their customers — expect to be able to make contact 24/7. Ensuring this happens requires good planning and UC system design to minimize the impact of disasters. This Report examines just what is involved in communications disaster planning, as well as the professional services opportunities it presents to partners.