

COMPANIES IMPROVE DATA PROTECTION AND MORE WITH COHESITY

APRIL 2017



We talked to six companies that have implemented Cohesity DataProtect and/or the Cohesity DataPlatform. When these companies evaluated Cohesity, their highest priority was reducing storage costs and improving data protection. To truly modernize their secondary storage infrastructure, they also recognized the importance of having a scalable, all-in-one solution that could both consolidate and better manage their entire secondary data environment.

Prior to implementing Cohesity, many of the companies we interviewed had significant challenges with the high cost of their secondary storage. Several factors contributed to the high costs including the need to license multiple products, inadequate storage reduction, the need for professional services and extensive training, difficulty scaling and maintaining systems and adding capacity to expensive primary storage for lower-performance services, such as group file shares.

In addition to lower storage costs, all the companies we talked to also wanted a better data protection solution. Many companies were struggling with slow backup speeds, insufficient recovery times and cumbersome data archival methods. Solution complexity and high operational overhead was also a major issue. To address these issues, companies wanted a unified data protection solution that offered better backup performance, instant data recovery, simplified management, and seamless cloud integration for long-term data retention.

Companies also wanted to improve overall secondary storage management and they shared a common goal of combining secondary storage workloads under one roof. Depending on their environment and their operational needs, their objectives outside of data protection included providing self-service access to copies of production data for on-demand environments (such as test/dev), using secondary storage for file services and leveraging indexing and advanced search and analytics to find out-of-place confidential data and ensure data compliance.

Cohesity customers found that the key to addressing these challenges and needs is Cohesity's Hyperconverged Secondary Storage. Cohesity is a pioneer of Hyperconverged Secondary Storage, a new category of secondary storage based on a webscale, distributed file system that scales linearly and provides global data deduplication and automatic indexing as well as advanced search and analytics and policy-based management of all secondary storage workloads. These capabilities combine to provide a single system that efficiently stores, manages, and understands all data copies and workflows residing in a secondary storage environment – whether the data is on-premises or in the cloud. There are no point products, therefore less complexity and lower licensing costs.

It's a compelling value proposition, and importantly, every company we talked to stated that Cohesity has met and exceeded their expectations and has helped them rapidly evolve their data protection and overall secondary data management. To learn about each customer's journey, we examined their business needs, their data center environment, their key challenges, the reasons they chose Cohesity, and the value they have derived. Read on to learn more about their experience.

MAJOR VALUE CUSTOMERS REALIZED WITH COHESITY

Cohesity customers we interviewed consistently told us their traditional storage systems weren't meeting their data protection and overall secondary data management needs. They said they needed simpler and faster data backup and recovery, more efficient and cost-effective data retention, better overall data visibility and self-service access to copies of production databases.

To achieve these data protection and secondary storage improvements, companies recognized they required a new secondary storage approach. They wanted a secondary storage and data protection solution architected from the ground up to be simple, scalable, and flexible. They found Cohesity DataProtect and the Cohesity DataPlatform to be the game changing solutions they were looking for. Cohesity DataProtect and the Cohesity DataPlatform leverage a hyperconverged secondary storage architecture that enabled them to both transform their data protection and consolidate secondary data workloads. Below are seven areas where companies we interviewed realized significant value with Cohesity.

- **Cost savings.** Cost savings was the number one reason customers switched to Cohesity. Customers realized savings in three areas: lower licensing costs, better data reduction ratios, and the ability to move file services off higher cost primary storage systems. Companies also realized significant operational savings with faster backups, substantial reduction in backup and recovery administration time and eliminating the cost and hassle of tape archival.
- **Simplified management.** Cohesity customers talked about the value of having an all-in-one solution that makes it easy to create data protection policies, quickly recover data, and intelligently migrate data to the cloud for long term data retention. Customers also cited Cohesity's automatic metadata indexing and the pattern matching that makes it easy to quickly locate specific files and find data that is non-compliant.

“I just can't say enough about the ease-of-use and the simple administration. What I like the most is we get the simplicity without sacrificing functionality.” (Ed Krupka, Burriss Logistics)

- **Easy to scale and maintain.** When they needed to expand their system, customers found Cohesity's hyperconverged architecture made growing the system easy, because each node in a cluster is self-sufficient with its own compute, memory, and storage. Customers also told us Cohesity's architecture simplifies system upgrades and recovery from node failures, because there's no disruption, no manual configuration, and no need for data migration.
- **Data efficiency.** Many customers found that Cohesity's global deduplication maximizes backup data reduction and they like the fact that separate deduplication appliances are not required, because it reduces solution complexity.

“We deduped 200TB down to 50TB, 25% of its original size. Cohesity offers a larger cluster backup, but we haven't needed it because of its incredible efficiency.” (Director of Security & Network Services, Law Firm)

- **Fast backups and recovery.** Multiple customers stated that Cohesity's backup speed is a real plus. Cohesity doesn't require weekly full backups and for critical backup workloads, they can leverage parallel data ingest and set policies that leverage Cohesity's SSD tier and optional post-process (vs. inline) deduplication to increase performance. Another feature customers highly value is rapid recovery. With Cohesity, companies can recover a failed VM from a clone of the backup VM and run the VM directly on the Cohesity platform while the data is still restoring. This capability provides fast access to critical data within minutes.

- **Cloud enablement.** Cloud-based storage is a central part of their data protection strategy. Some companies were using tape for data archival and with Cohesity they moved to simpler and more cost-effective data retention in the cloud.

“We saved a lot of money with Cohesity by archiving our data to Microsoft Azure and getting rid of tape that was expensive and very disruptive.” (Marlon Wenceslao, U. Pennsylvania)

- **Data workflow consolidation.** Lowering storage costs and/or improving data protection was the initial driver for all the companies we interviewed, but they quickly found Cohesity offers much more. Companies use Cohesity’s file services capability for NFS and SMB file shares and they also leverage Cohesity’s RESTful API to enable DBAs and other authorized users to quickly spin-up, refresh, restore and remove virtual copies of production data using tools and applications they are familiar with.

ARCHITECTURAL OVERVIEW

The foundation for Cohesity’s Hyperconverged Secondary Storage is their DataPlatform, a distributed, infinite scale-out file system that is targeted at secondary data workloads. The DataPlatform runs on Cohesity’s Hyperconverged Nodes, consisting of 2U rack mount boxes called blocks. Each 2U block can host up to four Cohesity Nodes, each of which has its own compute and storage. The storage component consists of both HDDs and Flash for high performance ingest and accelerated file services due to retention of frequently accessed blocks in Flash. Nodes are connected via dual 10Gb Ethernet ports. A minimum of three Nodes are required to form a Cohesity Cluster and clusters can be expanded to any size needed by adding one or more additional Nodes.

“We saw a tenfold improvement in speed with Cohesity. Part of that was due to switching from tape to disk, which saved us hours per week on backups. And each Cohesity node has network interfaces, which makes it considerably faster than other disk-based backup systems we looked at.” (Kris Kahl, Genex)

Nodes share nothing and there are no inherent bottlenecks in the distributed hardware architecture that Cohesity has implemented, so both capacity and performance scale linearly as Nodes are added. In addition, Cohesity’s hyperconverged architecture makes it much easier to grow the system as backup needs grow or to replace a node if there’s a failure. The administrator simply pops in another node, there’s no disruption, no need for data migration, and no manual configuration.

The base layer of the DataPlatform contains the file system, which is structured to be fully distributed and infinitely scalable. It consists of a set of services: a Cohesity-designed Cluster Manager that’s responsible for all Cohesity Cluster core services; an I/O engine that handles all cluster reads and writes; a file system metadata store, which is an indexing engine that scans and collects information about all data placed in a cluster; and a limitless snapshot facility based on Cohesity’s patented and highly efficient SnapTree technology.

“Cohesity is like the Google file system. It’s very easy to add nodes because it’s built to scale-out that way.” (Eirik Toft, XO Communications)

Layered on top of the file system is a storage services layer that provides interfaces for use by the secondary storage workflows such as Cohesity Analytics, Cohesity Test/Dev, Cohesity File Services and, most significantly, Cohesity DataProtect. This layer provides global, variable-length, user-configurable data deduplication; distributed SMB and NFS protocols; S3 object protocol; integrated Google-like global search capability; site-to-site replication; and AES 256-bit FIPS-compatible encryption with hardware acceleration. Finally, this layer is where Cohesity integrates with the public cloud - Amazon S3 and Glacier, Microsoft Azure, and Google Cloud Platform.

Cohesity's integrated cloud functionality consists of three policy-based capabilities. CloudArchive allows old snapshots to be moved from on-premises storage to the cloud for long term retention. CloudTier improves storage efficiency by moving cold data blocks into the cloud until its accessed again, at which point it's automatically moved back on-premises. CloudReplicate copies local snapshots to a cloud-based Cohesity Cluster for offsite data protection and disaster recovery in the cloud. In all three cases, data in the cloud is indexed for fast search and retrieval and it is encrypted both in-flight and at-rest.

The next layer up is the Application Layer, where higher level services needed by the integrated secondary storage workflows are provided. These are services such as scheduling, data archival, cloning, and policy management. This layer is also where the secondary storage workflows themselves reside – Analytics, Test/Dev, File Services and, most significantly, DataProtect -- a complete backup and recovery solution that is tightly integrated with the DataPlatform. DataProtect uses the globally de-duplicated, scale-out storage that eliminates the need for separate backup servers and software. It supports instantaneous RTOs and enables RPOs of less than 15 minutes and integrates seamlessly with VMware and all major cloud vendors, providing a substantial decrease in data protection costs and complexity.

A very interesting and important feature of the Cohesity platform also resides in this topmost Application layer -- the Analytics Workbench. This is an open application framework, including an SDK, that allows users and 3rd parties to create custom analytics modules and run them alongside the native Cohesity Analytics, greatly expanding the utility of the DataPlatform.

REPORTS FROM THE FIELD

In the following customer reports, six Cohesity customers share their secondary storage challenges, the reasons they selected Cohesity and the major benefits they realized when they moved to Cohesity hyperconverged secondary storage, which include lower storage costs, simplified management, better data protection performance, integrated cloud storage, and scalability.

Summary of Cohesity Customers Interviewed						
Company	XO Comm.	Law Firm	Univ. of Pennsylvania	Genex Services	Online Retailer	Burriss Logistics
Person Interviewed	Systems Engineer	Director Security & Network Services	Sr. Systems Mgr. of Systems & Infrastructure Services	Systems Mgr. / Systems Engineer	Systems Architect	CIO & President of Information Technology
Industry	Telecom	Legal	Education	Healthcare	Ecommerce	Distribution
Top Three Reasons Companies Made the Switch to Cohesity						
#1 Reason	Scalability	Global dedupe	Backup Speed	Simplified management	Scalability	Value
#2 Reason	Global dedupe	Simplified management	Cloud archival	Scalability	Performance	Simplicity & Support
#3 Reason	Simplified management	Better Data Protection	Scalability	Ease of upgrades	Lower Cost & Cloud	Performance

Ed Krupka – Burriss Logistics

Burriss Logistics is a privately-owned company founded in 1925. The family-owned firm has 17 locations and over 1,700 employees. Burriss provides custom supply chain solutions for controlled-temperature food distribution operations.

Ed Krupka is Burriss' Chief Information Officer and President of Information Technology.

WHY BURRIS CHOSE COHESITY

The primary data center in Delaware replicates to a data center in Philadelphia. Their legacy EMC environment had VMAX and VNX for production storage. RecoverPoint was used for replication and NetWorker for archival backups to Data Domains.

The Networker / Data Domain integration was complex and difficult. Ed said, "We paid for several service engagements to solve problems around Data Domain and Networker. We did not want to do that going forward."

Although Burriss has been an EMC shop for 20 years, integration problems and high costs encouraged Ed to look at innovative data protection platforms. He said, "The lights are always on somewhere in our environment and being down is not an option. That was the driver for us to go out on a limb and look at new storage platforms and data protection."

THE COHESITY ENVIRONMENT

Today the primary data center has about 250 host VMs and fewer than a dozen physical servers. Burriss replaced production EMC with Pure Storage and replaced Data Domain with seven Cohesity nodes at just over 150TB capacity. They have used about 77TB.

Cost savings and efficiency gains have been dramatic. The data center footprint has gone from 3 EMC cabinets to a 2U Cohesity rack, which lowers power and cooling costs. Burriss also eliminated the need for professional services and training. With every refresh, Burriss paid EMC \$250,000 for professional services plus additional expense for training. Cohesity's training and services charge has been zero, due to its much simpler tech stack and a dramatic improvement in ease of use.

Dedupe has also significantly improved available capacity over EMC. Data Domain's best dedupe ratio was about 33:1 ratio after three years of usage. Ed said, "We were pushing 100:1 with Cohesity. After ingesting the Postgres (PostgreSQL) database, dedupe ratios went to 46:1. That ratio is slowly increasing and we are more than pleased with it."

Data migrations to Cohesity went smoothly and was much easier than migrating archives with Data Domain. Ed said, "It was easier going from EMC to Cohesity than it was from EMC to EMC."

TOP BENEFITS

1. **Excellent value proposition.** Ed and his team wanted to reduce costs and increase value. Cohesity gave them an innovative data protection architecture at a much less expensive purchase price and far more operational value. Ed said, "The value is a major advantage for us. It's not just the purchase cost, although Cohesity is easily half of what it was costing us with EMC; it's also the simplicity and the lower cost of operations."
2. **Simplicity.** During this interview, Ed texted his infrastructure admin with a question about Cohesity. The admin texted back the exact dashboard view he needed. "It would've taken me a week to get that kind of figure on our old environment," Ed pointed out. Simple administration with good policy management lets them hold the line on employee headcount and upgrades have gone well with no downtime or maintenance windows. "I just can't say enough about the ease-of-use and the simple administration and what I like the most is we get the simplicity without sacrificing functionality."
3. **Better Performance:** Ed noted that "Backup speed is 2-4X faster than what we were seeing with EMC Data Domain and Networker."
4. **Support.** They have not had any emergencies with Cohesity. They contacted support with some initial configuration questions and issues and got the quick answers they needed. "It's been

great. Big vendors have a lot of capabilities and resources, but their support organizations are so hard to get through. I would give Cohesity support five out of five stars.”

WRAP-UP

“The team could not be more comfortable and pleased with the results and the performance in our environment. And Cohesity is passionate and they are doing some exciting things. We’re looking forward to hearing what’s coming next.”

Director of Security and Network Services – Large Law Firm

This large U.S. law firm has 10 remote offices throughout the country and employs 400 attorneys and 1100 total staff.

The Director of Security and Network Services oversees the firm’s main and backup data centers. Both data centers are 90% virtualized, and both have an 8-node Cohesity cluster.

WHY THE LAW FIRM CHOSE COHESITY

The firm initially used NetBackup in its Milwaukee data center. The legacy environment backed up Exchange, physical servers, and VMs to local storage. The Director preferred to keep backups for 90 days but there was not enough storage capacity. Another issue was the time it took to back up several 15 TB VMs. They also wanted to replicate to remote locations for improved data protection and they wanted an easy way to do backup verification and disaster recovery testing.

The Director and his team began looking for other options but many of them required installing agents on the VMs. They narrowed their shortlist to Cohesity and Rubrik. “Cohesity took half the time to perform backups, he said, “and Cohesity was straightforward and easier to use. I didn’t have to muck around. For me that was key. It was easy to use right out of the box.”

THE COHESITY ENVIRONMENT

The Director and his team installed 8-node Cohesity systems in Milwaukee and in the second data center in Phoenix. They backup 500 virtual servers in Milwaukee and replicate to Phoenix for data resiliency.

Global deduplication yields high storage capacity, which lets them easily store backup data for 90 days. Cohesity data retention policies automatically destroy older backups. Although Cohesity is highly scalable, the Director said, “Ninety days works for our business model. We have a very distinct data destruction policy.”

Data backup and restores are much easier now. In the legacy environment, administrators had to go to the backup application, login to the interface, search for the backup job, follow the folder structure, and hopefully find the file. They followed a similar process to locate snapshots. The Director said, “Cohesity backup scheduling is very straightforward and simple. We do a search and quickly find the files and choose the recovery point, it’s that easy.”

TOP BENEFITS

1. **Global dedupe.** Deduplication enabled the firm to crunch 200 TB down to 50TB, 25% of its original size. “Cohesity offers a larger cluster backup, but we haven’t needed it because of its incredible efficiency.”
2. **Simplified management.** Backup admins quickly search for files or snapshots and set restore points. “Backup Exec was very complex and difficult. It could take my admins all day to administer. Cohesity makes it easy to manage, and we’ve gone from all day to minutes. Also, before we had no good way to do DR testing. Now we do.”

- Data protection.** Cohesity covers all their backup in one hour at night. The Director and his team can easily verify backup from the dashboard and it's very simple to restore files or VMs. "After we got Cohesity, one of our servers failed big-time. It blew up. It took Cohesity just fifteen minutes to bring the failed machine back up and make it operational."

WRAP-UP

"We've seen very significant OpEx savings with simplified administration. With NetBackup, the big 15TB servers were taking three to four days! Now they take a few minutes. That's a huge difference."

Eirik Toft -- XO Communications

XO Communications, recently acquired by Verizon, owns and operates global IP and Ethernet networks that provides secure private data networking, cloud connectivity, unified communications, voice, Internet access, and managed services.

Eirik Toft is a Systems Engineer working with a large XO client. This major government entity stores massive amounts of data, and requires robust data protection with geographic redundancy. Eirik and his team deployed two Cohesity data center environments on opposite coasts.

WHY XO CHOSE COHESITY

When XO won the government entity's business, they needed to build multiple data centers capable of supporting intensive applications and large data volumes. One of the critical elements was a highly scalable storage system that could efficiently store and manage massive amounts of data. They owned EMC storage and were interested in Isilon, but its high cost was a serious drawback.

One of Eirik's peers suggested Cohesity and its innovative hyperconverged architecture. Eirik contacted them and they offered what he needed at less than half the price of the other quote. He said, "We had an immense storage requirement and cost was a big driver. Cohesity's cost per terabyte of storage was fantastic."

Deployment was as simple as mounting an NFS server. Cohesity throttled the initial full backup so it did not impact network processing. Now an incremental backup of over 200 VMs completes in less than 10 minutes, and they have done six upgrades with no downtime.

After the sale, Cohesity frequently introduced new features. Eirik said, "They have capabilities that we didn't even know we needed. And when we realized they had them and saw how useful they were, we adopted them. This just validates our purchase even more."

THE COHESITY ENVIRONMENT

Today's primary environment runs on HP DL 580 hosts that is virtualized with VMware vSphere 5.5. Cohesity mounts as a share and backs up about 200-250 VMs at each site. The virtual platforms combine Windows and Linux hosts. There are also several application testing labs using VMs and Cohesity for easy spin-up and destruction. Cohesity automatically assigns correct backup policies to new VMs.

The team uses Cohesity native file search storage for fast and efficient searches. They also appreciate its stand up and tear down abilities for dynamic VMs. Eirik said, "It's kind of fun to watch. Once you remove the VM from a test environment, you'll see Cohesity tear down the NFS data store from the host. It cleans up after itself."

Eirik also uses Cohesity's scripting support. "I'm writing custom scripts using their API. Implementing extra code into their system is very easy and very forgiving, and it points out where you made a mistake. It makes customizing very easy."

TOP BENEFITS

1. **Easy scalability.** “EMC is monolithic and I have to add racks to scale. Cohesity is different. It’s very easy to add nodes because it’s built to scale-out that way.”
2. **Deduplication.** Dedupe is critical for storing data. In the test environment, Eirik realizes a 5.75:1 ratio, and 4.74 TB of raw physical storage presents almost 250 TB of logical storage. The production server backup is largely logs, zipped files, and SQL databases. Even though this data does not have much repetitive content, XO still sees a dedupe rate of 2.2:1.
3. **Policies.** Policy-based scheduling is highly efficient and saves on management time. Eirik remarked, “The retention policies can be as simple or as complex as you need them, which is very beneficial. I like that there don’t seem to be any limits in the policies. At least I haven’t hit any yet.”
4. **Support.** Support is skilled and responsive when Eirik needs them, which isn’t often. “We had one issue with Cohesity and we didn’t even know we had an issue. Cohesity told us about it and came out to fix it.” Cohesity is extremely open to suggestions and requests for product enhancements.

WRAP-UP

“I tested the Cohesity environment by doing everything I could think of to break it: pulling cables, turning the machines on and off. Short of hitting a node with a hammer, there was nothing I could do to destroy it or screw it up.”

Online Retailer

This online retailer is one of the fastest growing ecommerce sites in the United States. Two data centers are located on the east coast and several fulfillment centers are spread across the country. We spoke with the System Architect responsible for the company’s virtualized infrastructure including compute and storage.

WHY THE COMPANY CHOSE COHESITY

The company used a major storage vendor for data protection and archival. The secondary storage system was slow and running out of capacity, so the System Architect and his team researched new options for a technology refresh. They selected Cohesity for its high scalability, low overall cost, and ability to consolidate secondary storage workloads under one roof.

The System Architect said, “With Cohesity, you consume only what you need now, and it’s a simple matter to add more as you go. Instead of refreshing the other system every two years and paying for surplus capacity upfront, with Cohesity we just add nodes as needed on a quarterly basis.”

THE COHESITY ENVIRONMENT

The System Architect and his team deployed two Cohesity 4-cluster nodes, one in each data center. They will be expanding soon. The data centers run both Windows and Linux and are 100% virtualized with VMware. About 50 VMware hosts reside in the data centers for production and test/dev, and another 7 are in the fulfillment centers.

The clusters are protecting about 40TB of primary data in the main data center, and about 12TB in the secondary test/dev data center. Cohesity averages a deduplication ratio of 50:1, a significant improvement over the 10:1 they were getting with the legacy storage system.

Cohesity enables the team to remotely backup from the fulfillment centers to a central repository. The System Architect said, “Remote backup is something else that’s great about Cohesity. Backups

have gone from hours to minutes.” They are using cloud-based archival for cold storage, where Cohesity writes to Iron Mountain’s Cloud.

Data protection was their initial driver, but the team found that Cohesity offers much more. The System Architect said, “Data protection was our big use case when we adopted Cohesity. We then began exploring its capabilities, and saw interesting new use cases.” The team uses Cohesity’s file services capability to expose NFS and SMB file shares, including an NFS mount for the Frozen Splunk log data. They are also working closely with Cohesity and VMware to create an “easy button” for DBAs, which leverages Cohesity’s RESTful API to automatically refresh and restore replicated virtual copies of their Microsoft SQL database.

TOP BENEFITS

1. **High scalability.** Their legacy storage could only scale so far before the limited compute of the controller becomes a constraint. “This is where Cohesity has such a good value proposition. You have a cluster that can easily scale out to one thousand nodes because each node has its own compute, memory and storage.”
2. **Backup performance.** Admins can adjust policies to optimize workloads. For example, a policy leverages the Cohesity SSD tier and post-process (vs. inline) deduplication to increase data ingest performance. “Cohesity has a lot of flexibility and backups are very fast. Before, our Veeam Change Block Tracking (CBT) would take over an hour. With Cohesity it takes less than five minutes.”
3. **Low costs.** The company paid separate licensing and maintenance costs for storage hardware and software. “With Cohesity, we have one vendor and one competitively priced solution that includes hardware and software. That has made a big difference in cost.”
4. **Cloud enablement.** The team uses Cohesity for cloud-based archival and writes to Iron Mountain’s Cloud using a certified S3 connector. They recently tested a cloud data restore on a virtual application with 200 GBs of archived storage. “We only needed to transfer 5% of data to do a full recovery. This a big plus, because we avoid paying the high data transfer costs you often see with cloud storage.”

WRAP-UP

“Every time you add storage to Cohesity, you also add compute. Instead of running slower and slower and slower, it runs faster and faster and faster.”

Marlon Wenceslao – Annenberg School for Communication at the University of Pennsylvania

Annenberg is one of twelve schools at the University of Pennsylvania. The University maintains a central IT department and each school also has its own LSP (Local Support Provider). Annenberg’s graduate students generate large amounts of communications research, and projects have high compute and data storage requirements.

Marlon Wenceslao is Senior Systems Manager of Systems & Infrastructure Services. Cohesity hyperconverged secondary storage allows the Infrastructure Team to easily scale storage without expanding VCE Vblock primary storage.

WHY ANNENBERG CHOSE COHESITY

The legacy environment ran Veeam, Backup Exec, and EqualLogic Dell storage. Veeam still had two years on its licensing contract, so Marlon planned to replace the aging physical backup storage infrastructure.

Marlon first looked seriously at Dell and Exablox. He was particularly interested in Exablox, but its lack of hyperconvergence and large footprint put him off. Then he looked at Cohesity and Rubrik. He liked what he saw with Cohesity, they had stronger protocol support including support for both NFS and SMB, so he bought it to replace the legacy Dell storage.

THE COHESITY ENVIRONMENT

Once Veeam was backing up the virtual environment using Cohesity as a Repository, Marlon agreed to test Cohesity's native VM backup capability. The legacy environment took almost two days to finish full virtual backup. Cohesity took three hours. He said, "I was surprised at how fast it was. It was easy, simple, and just what I was looking for. Even though it's incremental, it's fully hydrated. I didn't expect Cohesity to do that." Marlon moved 80% of backups from Veeam to Cohesity. Later he phased out Backup Exec in favor of Cohesity for physical server backup.

Annenberg uses Vblock Systems to store high performance primary data and file shares, but scaling is expensive and data center space is limited. Instead of growing the Vblock footprint, Marlon decided to host file shares on Cohesity. He said, "We might as well use secondary storage to host the file shares. It's a big footprint advantage with Cohesity: we don't have to scale Vblock, and Cohesity dedupe/compression is huge. Cohesity also provides a redundant architecture, so if a node goes down the data is still available. We didn't have that with VM hosted shares on Vblock systems."

Security compliance is another important benefit. Annenberg does not allow students to save data containing personally identifiable information (PII). Admins use Cohesity Pattern Finder to locate non-compliant text patterns like Social Security numbers.

Marlon also likes dynamically creating and destroying VM clones. For example, if his web admins need to upgrade Drupal or WordPress, Marlon easily creates a cloned environment for upgrade and validation. Once he gets the green light, they update the live environment and destroys the clone. "It's a very easy two-step process. Clone the VM, then destroy the VM. That's it."

TOP BENEFITS

1. **Backup speed.** Since Cohesity snapshots take under a second, the team can take many snapshots without impacting backup performance. This enables them to create numerous restore points. Marlon said, "The legacy environment took two days to finish a full virtual backup. Cohesity took only three hours."
2. **Cloud archival.** Annenberg keeps archival data for five years. They are moving forward with a cloud-first strategy for archives, which replaces storing long-term archival to tape. Marlon said, "With Cohesity we saved a lot of money by archiving our data to Azure and getting rid of tape that was expensive and very disruptive." Annenberg also plans to use Cohesity for cloud-based DR in the coming months.
3. **Scalability.** Marlon appreciates Cohesity's hyperconverged secondary storage with its small footprint, backup speed, and efficient scalability. "It's a webscale kind of infrastructure where you don't have to worry about adding storage or compute or disk drives. It's just a matter of adding a node or a box, and you're good to go."
4. **Ease of management.** Marlon's small team needed to cut down their management overhead. Cohesity's policy-based management is key. "Cohesity policies are very helpful and the policies are simpler than Veeam. I can just use my iPhone to connect to my Cohesity cluster anywhere to update a policy or schedule a backup. It's got great flexibility and a rich toolset."

WRAP-UP

"I'm really pleased with Cohesity, how easily it does backup and how fast it does restore. Here's a case in point. If I had to restore an older VM it would take me at least five days to recall the tapes, catalog the tape, find the data, and restore VM. With Cohesity it takes just seconds. It's truly disruptive."

Kris Kahl and Mathew Varughese – Genex Services

Genex Services was founded in 1978 as a case management service for injured employees. The company focuses on compassionate case management, which reduces health care and disability expenses for their customers. Headquartered in Wayne, PA, Genex operates across the U.S. and Canada.

Mathew Varughese is Senior Systems Engineer and Kris Kahl is Systems Engineer charged with backup administration.

WHY GENEX SERVICES CHOSE COHESITY

Thirteen branch offices use Symantec NetBackup to backup their servers to the corporate data center in Wayne. Corporate IT used NetBackup to write local and incoming backup to tape, and then shipped it off-site for long-term data retention. The remote to data center backup worked well enough, but the team found that backing up a growing amount of data to tape took far too long. They decided to find a high-performance backup for the data center.

They considered retaining NetBackup and backup to disk, but Symantec licensing and maintenance was expensive. It was also difficult to scale and to add new backup targets. The team decided to look for new solutions that backed up both virtual and physical servers with better scalability, higher performance, and simplified management. They found Cohesity, which gave them the features they were looking for at a significantly lower price point.

THE COHESITY ENVIRONMENT

They initially invested in a proof of concept Cohesity system with four nodes. The test went well, so they added an additional four nodes and purchased a duplicate eight node system for their backup site. Data replicates between the two sites.

Cohesity was basically plug-and-play. Kris said, “One time a hardware node failed. The remaining nodes took over processing and Cohesity sent a new node. We plugged it in, the Cohesity engineers checked it out to make sure it was good, and everyone was happy.”

They soon saw several operational advantages. NetBackup would do incremental backup daily, plus a full weekly backup. But one server alone stores nearly 2 TB, and the change rate on it is only 50 to 60 MB a day. A full backup would take hours even with such a small change rate. With Cohesity, once it finished an initial full backup every subsequent backup is differential only and fully recoverable. So instead of taking several hours once a week to backup that server, it takes a few minutes every day.

Kris said, “We saw a tenfold improvement in speed with Cohesity. Part of that was due to switching from tape to disk, which saved us hours per week on backups. And each Cohesity node has network interfaces, which makes it considerably faster than other disk-based backup systems we looked at.”

Another feature that the team appreciates is that with Cohesity, you can recover a failed VM and start using it while Cohesity is still transferring data back to the ESXi data store. Kris said, “When you have a critical application, you can get access to it within a few minutes while the data is still restoring. This is very attractive.”

TOP BENEFITS

1. **Simplified management.** The single backup management pane saves backup administrators energy and time. “One of the big reasons that we went with Cohesity was its all-in-one single pane application.”
2. **Easy scalability.** System expansion doesn’t happen every day, but when it happens the process is simple. “Cohesity is very easy to expand, much easier than the other three backup products that I’ve used in the past.”

- Ease of upgrades.** When Cohesity upgrades with patches or OS upgrades, the process is dynamic and problem-free. “We’ve upgraded the clusters several times with patches and new features. Cohesity engineers can do it remotely, or we can do it ourselves. You basically hit a button and walk away. There’s no downtime, so the ease of upgrade is great.”

WRAP-UP

“Cohesity is pretty much all-in-one. It’s all accessible via a webpage, so there isn’t an agent or an application that you have to install on your laptop or PC. You access it from anywhere.”

TANEJA GROUP OPINION

For most enterprises, evolving data protection and achieving better economics is a critical part of their overall IT modernization agenda. This is indeed the case for all the companies we interviewed for this report. Cohesity customers all wanted lower storage costs, simpler and faster data backup and recovery and more efficient and cost-effective data retention.

But companies said they couldn’t make progress because they were strapped by system complexity, high licensing costs and the inability to easily grow their legacy systems. To achieve their objectives, organizations soon recognized they needed to look beyond the boundaries of their legacy storage systems and standalone data protection applications. This led them to Cohesity and they soon discovered the power of hyperconverged secondary storage. They found that it provided the ease of use, global data reduction and scalability needed to truly lower their storage costs and simplify and speed up their data protection processes as well as incrementally expand their secondary storage as needed without affecting performance.

Cohesity customers also had a strong desire to continually advance their overall secondary storage environment. They wanted better overall data visibility, self-service access to copies of production databases, lower cost file services and a clear path to the cloud. Once again, Cohesity proved up to the task by providing flexible policy-based orchestration for centralized management of multiple data workflows and the automation of metadata indexing to enable Google-like searching as well as advanced reporting and analytics to assist with planning and data compliance. Companies appreciated how Cohesity seamlessly incorporates cloud-based data workflows for archival and disaster recovery, so they can benefit from the convenience of on-demand cloud storage while still maintaining complete visibility, security, and control over their data irrespective of where it resides.

The term “paradigm shift” is often overused, but for hyperconverged secondary storage we think the term applies. There’s no doubt that Cohesity’s hyperconverged secondary storage approach has fundamentally changed data protection and secondary storage management for Cohesity customers. Cohesity’s simplicity, scalability and flexibility are the antithesis of what many companies are used to. And given the tremendous value and major benefits Cohesity customers are consistently realizing, we believe hyperconverged secondary storage will play a very strong role in shaping the future of data protection and secondary data management for the companies we interviewed as well as many more organizations for years to come.

NOTICE: The information and product recommendations made by Taneja Group are based upon public information and sources and may also include personal opinions both of Taneja Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Taneja Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.