COHESITY DATA PLATFORM

Hyperconverged Secondary Storage

MARCH 2016

Primary storage is often defined as storage hosting mission-critical applications with tight SLAs, requiring high performance. Secondary storage is where everything else typically ends up and, unfortunately, data stored there tends to accumulate without much oversight. Most of the improvements within the overall storage space, most recently driven by the move to hyperconverged infrastructure, have flowed into primary storage. By shifting the focus from individual hardware components to commoditized, clustered and virtualized storage, hyperconvergence has provided a highly-available virtual platform to run applications on, which has allowed IT to shift their focus from managing individual hardware components and onto running business applications, increasing productivity and reducing costs.

Companies adopting this new class of products certainly enjoyed the benefits, but were still nagged by a set of problems that it didn't address in a complete fashion. On the secondary storage side of things, they were still left dealing with too many separate use cases with their own point solutions. This led to too many products to manage, too much duplication and too much waste. In truth, many hyperconvergence vendors have done a reasonable job at addressing primary storage use cases, , on their platforms, but there's still more to be done there and more secondary storage use cases to address.

Now, however, a new category of storage has emerged. Hyperconverged Secondary Storage brings the same sort of distributed, scale-out file system to secondary storage that hyperconvergence brought to primary storage. But, given the disparate use cases that are embedded in secondary storage and the massive amount of data that resides there, it's an equally big problem to solve and it had to go further than just abstracting and scaling the underlying physical storage devices. True Hyperconverged Secondary Storage also integrates the key secondary storage workflows - Data Protection, DR, Analytics and Test/Dev - as well as providing global deduplication for overall file storage efficiency, file indexing and searching services for more efficient storage management and hooks into the cloud for efficient archiving.

Cohesity has taken this challenge head-on.

Before delving into the Cohesity Data Platform, the subject of this profile and one of the pioneering offerings in this new category, we'll take a quick look at the state of secondary storage today and note how current products haven't completely addressed these existing secondary storage problems, creating an opening for new competitors to step in.

CURRENT ISSUES IN SECONDARY STORAGE

Whereas primary storage architectures are moving towards unified, scale out, distributed file systems, legacy secondary storage has been mostly stagnant and has remained mired in three different issues:

1) Most of the data residing in secondary storage is "dark data" and therefore of little use

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

We estimate that roughly 80% of all customer data is in secondary storage, with a growth rate around 50% per year. Customers typically can't explain why it's growing so fast, how many duplicate copies there are or what, exactly, is even in it at all. This means that the majority of customer data is just sitting in secondary storage, opaque and useless to the business, representing a significant cost with no corresponding benefit without a lot of additional management and processing. The numbers show that it's an out-of-control situation.

2) It's fragmented

Secondary storage is the catch-all for things such as data protection, DR, test/dev, analytics, archiving (including to the Cloud), etc. Each of these domains is served by its own ecosystem of vendors providing their own software and hardware to address application-specific issues. When you start considering deduplication, Hadoop analytics, copy data management and the rest, it's not exaggerating to say that there are 10's of different systems and software packages involved. This creates a myriad of user interfaces to learn; applications to license, train on and maintain; and silos of specialized hardware. In the current world of legacy secondary storage, it's all necessary, but it forms a collection of expensive pain points.

3) It's rife with inefficiencies

Currently, data protection typically consumes a third of secondary storage volumes, and has the greatest amount of data duplication and over-provisioning. Backups placed in secondary storage are an insurance policy, generally ignored until something goes wrong and a restore needs to be done. This consumes a lot of space without providing much corresponding value beyond peace of mind.

"Dark data" often represents another third of total secondary storage. When someone finally does decide to peer into the "dark data" in secondary storage, it's usually copied out to another location, such as a Hadoop cluster, for data mining. This requires another stack of hardware and an additional copy of the data (more storage space consumed), plus network bandwidth to move the data.

Lastly, test-dev instances are yet another large consumer of secondary storage. If code needs to be updated, or a bug is found and needs troubleshooting, the engineering team must have access to a suitable code/debug environment to address the issue. Since secondary storage is generally just that, a passive pool of storage, such test/dev activity usually requires yet another hardware stack and an additional copy of data in order to create a real world environment to code on or debug against.

Generally, a copy is required for every specialized usage of data in secondary storage. This often requires translations from one format to another (restoring a backup from tape to a disk based file system, for example). This activity may be necessary to make the data usable for a particular application, but it's a pure overhead cost. Then, as "collateral damage", this proliferation of multiple copies of the same or similar data creates a large copy data management (CDM) problem. Even with expensive deduplication appliances, data proliferation is still an issue; copies may be smaller, but they still need to be managed.

WHY SHOULD I CARE?

The question becomes, "Why does any of this matter?" The most obvious issue is cost – all of the extra storage equipment, specialized software, licensing fees, trained personnel, etc. are expensive. Implementing a simpler, more effective secondary storage architecture can eliminate most or all of these excess expenses and free up people to work on more business critical tasks.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

It's also a matter of efficient asset utilization. Data is a business asset, and an exceptionally important one at that. With 80% of customer data residing in secondary storage, while being the fastest growing data segment, that's a huge asset to allow to go "dark" and become stale. A secondary storage architecture that allows you to easily crack open and mine all of that data would be a significant improvement over the situation that exists today. It's also dangerous to not know what you have and where you have it. You don't want to lose track of sensitive information or data that's governed under various regulatory regimes. That can create security issues or put your business at risk of being fined for regulatory non-compliance.

There's also a fundamental mismatch between the legacy secondary storage architectures in use today and some of the recent innovations in primary storage that are being driven by the adoption of hyperconverged infrastructure. Hyperconvergence vendors are delivering distributed, infinitely scalable file systems that are tailored for highly flexible, virtualized workloads and that greatly simplify system management and expansion. Moving to such a system and then grafting a tape library or data deduplication appliance onto it for backup and setting up a second system in parallel for test/dev or data mining negates many of the improvements gained by moving to a hyperconverged system in the first place.

CURRENT POINT SOLUTIONS FALL SHORT

Most storage architecture improvements have been focused on primary storage, but in recent years, there have been a few piecemeal attempts to improve the situation in secondary storage. Copy Data Management (CDM) reduces the amount of clutter in storage by minimizing the number of full copies that are saved, but doesn't do much about fragmentation or dark data. Data deduplication technologies reduce the amount of redundant data that's stored, but only on a local device basis. Newer, cloud-based data protection architectures, where only a few recent backups are kept locally, with the rest automatically archived in the cloud, do help keep local storage consumption down, but merely move the problem somewhere else.

Fundamentally, all of these efforts fail at solving the bigger problem because they're point solutions that don't address it holistically. At the end of the day, you're still left with:

- Siloed data
- Multiple copies of data to deal with
- Multiple media servers
- Multiple deduplication devices, inefficient deduplication
- No global QoS capabilities governing the whole system

Given this list, it wouldn't be uncommon at all to have 20 or more products from 5 or more vendors dedicated to secondary storage operations when you're throwing point solutions, no matter how good they are, at the space.

ASSESSING THE COHESITY DATA PLATFORM ARCHITECTURE

As stated earlier, any platform intending to provide what we view as true Hyperconverged Secondary Storage is going to need to provide a web-scale, distributed, infinitely scale-out file system, integrate the key secondary storage workflows, offer global deduplication and provide file indexing and searching services.

The consolidation of all secondary storage use cases onto a single storage space means that the platform must be able to handle multiple, concurrent, mixed workloads. There are three critical capabilities that are needed for this:

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

1) Perform data operations extremely well – the system must be able to provide sustained throughput for sequential I/O at the same time that it's delivering high IOPS for random I/O

2) Perform metadata operations very efficiently – managing many small files, creating and deleting files, dealing with directory structures

3) Provide performance isolation – i.e. one huge backup job shouldn't starve other workloads

With this in mind, let's take a look at the Cohesity Data Platform Architecture and see how it does.

The Cohesity Data Platform is a scale-out platform targeted at secondary storage. It deduplicates and indexes all data upon ingestion for maximum efficiency and also integrates the major secondary storage workflows - data protection, test/dev, analytics and file services – as part of the platform. It's built on a four layer, web-scale architecture that has, at its heart, a distributed, infinite scale-out file system.

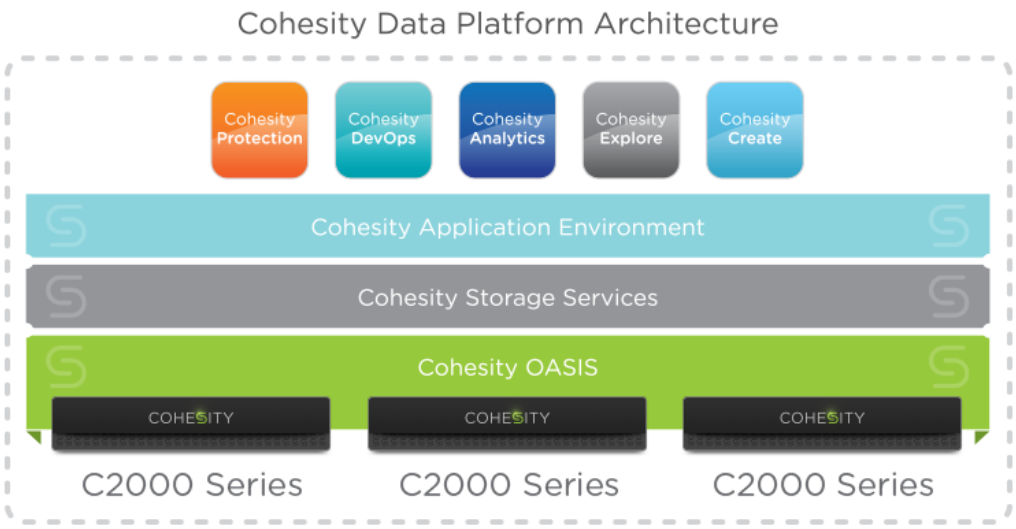

Layer 1 – The Physical Layer

The bottom, physical layer, consists of low cost, high performance commodity hardware, packaged in 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. Each Node consists of Flash and HDDs for high performance ingest, as well as for 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. 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.

Layer 2 – Base Layer: Cohesity OASIS File System

The OASIS (Open Architecture for Scalable Intelligent Storage) file system is the heart of the Cohesity Data Platform. It consists of a set of services:

- The Cluster Manager is responsible for all Cohesity Cluster core services, such as configuration management and monitoring component status. Instead of using one of the available open source cluster managers, Cohesity custom-built their own to provide the high levels of fault tolerance and performance that they deemed necessary.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

- The I/O engine handles all cluster reads and writes. In addition to *jointly optimizing both sequential and random I/O streams via automatic tiering between the HDD and SSD layers (Critical Capability #1)*, it also concurrently writes data to two nodes in the cluster for fault tolerance.

- The Metadata Store, as implied, holds the file system metadata and is *highly optimized for metadata retrieval (Critical Capability #2)*. Also of note, it ensures that there are always three copies spread around the cluster so that the metadata is well protected from failure.

- The Indexing Engine performs a two-pass inspection of all data that's placed in a cluster. On the first pass, it collects high level object information, such as the names of VMs, for quick object retrieval. On the second pass, it peers inside the high level objects to collect content information, such as filenames. This enables native search and recovery for both high level objects and their contents.

- The Integrated Data Protection Engine operates in conjunction with 3rd party services, such as VMware's VADP interface, to provide customers with complete data protection native to the platform.

OASIS is structured to be infinitely scalable and fully distributed, as are other hyperconverged file systems. However, it has several features that make it unique. First off, it was designed with a flexible, open architecture to allow Cohesity to achieve their vision of consolidating multiple secondary storage workloads onto a single, coherent platform.

Second, it was built to be strongly consistent. Operations on distributed systems are not atomic, meaning that they can complete at different times on different nodes, leading to potential data corruption. The traditional workaround for this was to assume that data written to a distributed system would *not* be consistent on all nodes at the time of the write, but that eventually all nodes would become consistent as the write operation cascaded through the system. This concept of 'Eventual Consistency' is fine in most cases, but for the level of performance that Cohesity was looking for it wasn't good enough. They wanted any piece of data that was just written to the cluster to be immediately available to be read right back in the next operation, so to deliver the consistency and performance that were needed, Cohesity built a custom NoSQL store for enhanced horizontal scaling across the cluster and used Paxos consensus-resolving protocols to avoid the Eventual Consistency problem. The result is a high performance, strongly consistent, highly scalable file system.

A final key feature of OASIS is the intelligent placement of data within the cluster. As data is read into the file system, a copy is immediately written to another node in the cluster to provide fault tolerance. If a customer has more than one block, Cohesity will place this second copy on a node in a different block to increase the level of fault tolerance even more. Customers wanting the highest level of fault tolerance can choose to have more than two copies created and spread around the cluster. As part of this process of writing and replicating data to a cluster for fault tolerance, Cohesity also distributes it evenly across all nodes in the cluster to avoid creating any I/O bottlenecks, maximizing performance in addition to fault tolerance.

Layer 3 – Cohesity Storage Services

This is where the higher level interfaces that present the OASIS file system to the storage workflow applications reside. In addition to providing services to the application layer, it allows the user to begin to make use of the powerful capabilities provided by the OASIS file system. Building on top of the Indexing Engine, it provides an integrated Google-like search box capability for files and objects stored in OASIS. Access to a native MapReduce integration is also available in this layer.

One essential service residing in this layer, snapshots, is based on Cohesity's patented SnapTree technology. Traditional file system snapshots use a linear linked chain technique, where the chain

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

grows longer as snapped changes are added to it. When it's time to retrieve data, the system has to work back through the links to recover the desired version, meaning that the time to reconstruct a data object increases as the length of the data chain increases. With SnapTree, Cohesity uses a tree structure instead of a chain to hold pointers to the data blocks. This allows them to limit the number of pointers that they have to de-reference to access data blocks to exactly three, no matter how many snapshots back they have to go to access the data they want. In addition to ensuring all data copies remain fully hydrated, it allows users to take an infinite number of snapshots at any interval of time and keep them forever.

The Services layer is also where the Cohesity Data Platform provides a foundational site-to-site replication capability, to support disaster recovery and business continuity. The user can set up replications to a different volume on the same cluster, to a separate on-premises or offsite cluster or to one of the public cloud options that they're integrated with. Additionally, replication granularity is also user selectable – individual backup jobs or an entire logical volume can be replicated. For an even higher level of protection, all data on one cluster can be replicated to one or more other clusters.

This Services layer also supports the SMB 3.0 and NFS storage protocols, allowing the CDP to be used as a NAS target for simpler file sharing. 3rd party data protection and other applications can also utilize these protocols to seamlessly use it as a storage target.

Data deduplication is a must-have technology in any storage system for eliminating redundant data blocks and maximizing storage capacity. Instead of the more common file or block level deduplication that only works across a single storage pool, Cohesity chose to implement a type of variable length data deduplication technology that works globally across an entire cluster for maximum storage efficiency. They also gave individual users control over it, allowing them to choose to turn it off completely, or having it run either in-line or post-process, depending on the characteristics of their specific workload.

Compliance with federal regulations for data security in the healthcare and financial services industries requires strong encryption. Cohesity has implemented AES 256-bit FIPS-compatible encryption, with hardware acceleration, for this purpose.

Finally, this layer is where Cohesity's integration with the public cloud occurs, notably with Amazon S3 and Glacier, Microsoft Azure and Google Cloud Storage Nearline. Users can take advantage of this capability to seamlessly set up policy-based offsite archival for their data. Using what Cohesity terms "Cloud Tiering", users can further utilize S3 as a temporary storage capacity extension to a cluster, effectively renting extra cluster space in the cloud as needed.

Layer 4 – Cohesity Application Layer

This layer provides the high level services needed by all of the secondary storage workflows that Cohesity is consolidating, such as policy management, scheduling, cloning and data archival. Additionally, it hosts Cohesity Protection (initially integrated with vCenter, but targeting other environments for the future), Cohesity Analytics (consisting of modules for monitoring storage utilization trending, reporting on user, VM and file data, log filtering and virus fingerprint scanning), Cohesity Test/Dev and Cohesity File Services.

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

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

BRINGING IT ALL TOGETHER

Clean, well-designed architectures, such as one underpinning the Cohesity Data Platform, are technically satisfying, but they still have to prove their value in the real world. Since the main user touchpoints are the integrated workflows enabled by their architecture, the only test that matters is how well Cohesity delivers on them. Here's how they do:

Data Protection: The Cohesity Protection workflow, fully integrated with VMware vCenter and using VMware's VADP (vSphere APIs for Data Protection), is a complete, agentless data protection package for VMware environments. Using the OASIS Indexing Engine, which operates cluster-wide and scales out accordingly, customers can index their entire vCenter contents to quickly get a high level view of their virtual environment and begin protecting their VMs. Then, as VMs are protected, the Indexing Engine will peer inside them and index their metadata, giving it a complete view of all of the VMs running on the cluster and the contents of their individual file systems, and enabling restores of either complete VMs or individual files. Due to Cohesity's built-in SnapTree snapshot technology, customers can also afford to take large numbers of time and space efficient snapshots, which translates into very aggressive RTOs and RPOs. Specific RPO/RTO numbers depend on several factors (how many VMs you're trying to protect at once, how much compute power you've got, …), but you can take snapshots at sub-minute intervals if you wish.

This combination of comprehensive knowledge of the virtual environment, along with very granular, essentially infinite snapshots, delivers a solid backup capability. However, the most thorough backup process isn't really worth much if restores aren't quick and easy. Cohesity takes care of this by simply taking a snapshot of a desired VM restore point and powering it up directly on the Cohesity platform itself for virtually instantaneous recovery. Once powered up, it can remain running on the Cohesity box or it can be moved elsewhere using Storage vMotion.

While the core replication capability resides in the Services layer of the Cohesity architecture, the Cohesity Protection module allows users to set up higher level DR functions, such as automatic failover to a clone/set of clones running on a separate cluster.

Even though they've done the work to provide thorough end-to-end native data protection for VMware environments, Cohesity also realized that certain customers in certain situations might want the ability to use 3rd party data protection products, so they've enabled them to use the Cohesity Data Platform as a backup target via their NFS and SMB 3.0 interfaces. This allows customers to still benefit from the secondary storage consolidation offered by Cohesity, while maintaining flexibility to use the data protection software of their choice.

Analytics: As with Cohesity Protection, Cohesity Analytics also leverages the OASIS Indexing Engine to look at all data on the cluster. Out of the box, this yields high level information concerning storage utilization and available capacity, file type and user access history and VM storage consumption, things that are all useful for trend analysis and capacity planning.

Cohesity's integrated MapReduce framework allows large datasets to be processed right on the cluster, instead of having to move them off to a specialized compute stack. This not only reinforces the Cohesity vision of eliminating sprawl and consolidating secondary storage, it provides a very powerful method for lighting up the dark data that resides there.

Finally, the Analytics WorkBench allows users to extend the native analytics capabilities of the platform by creating custom analytics for their specific business needs. For the first time, customers will be able to customize and utilize their secondary storage space as a productive business tool instead of as a passive data area. It will also enable an entire ecosystem to spring up for providing applications that run on the Cohesity Data Platform.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Test/Dev: Typically, when developers needed to update a code base or track down a defect, IT would set up a separate compute/storage stack for them to work on, creating yet more sprawl and adding another set of equipment to manage. With Cohesity Test/Dev, the need to do that goes away. Developers can take a snapshot of any backup of their production system and clone it to run directly on the Cohesity Data Platform to create their own development and testing environments. The automated cloning of backups is also available to spin up zero space clones even more quickly, streamlining the test/dev process even further. Using existing idle backups for this test/dev work also helps eliminate the propagation of excess copies of redundant data, a key Cohesity objective.

File Services: Cohesity File Services allows users to access their files stored in the cluster over NFS or SMB. Administrators, utilizing the Cohesity single pane management interface, can then track and analyze usage patterns. Administrators can also create file shares that are individually optimized for specific workloads. Users can specify whether their workload is either CPU intensive, or disk I/O intensive, and can also specify how they want to apply deduplication and whether or not encryption and compression are turned on.

Cohesity has implemented their *performance isolation and QoS management capabilities (Critical Capability #3)* as a policy-based mechanism that cuts across Cohesity Protection and Cohesity File Services. Users are able to characterize their workload as either a test/dev job (random I/O) or a backup job (sequential I/O). System Flash storage is prioritized for use with the test/dev workloads, while backup workloads are more focused on the HDDs. For even more granularity, either job type can also be ranked as 'High' or 'Low' priority. These dual type/priority settings are then mapped into a proportional resource allocation mechanism embedded in the system to enable it to make appropriate resource trade-offs among the various workloads.

Additionally, Cohesity has implemented user-enabled adaptive throttling for backup workloads. When activated, it monitors the HDD latencies and dials back the data intake rate if it hits a user-defined latency threshold. This feature, in conjunction with the job type/priority QoS settings and the priority-dependant internal resource allocation mechanism gives the user a high degree of control over how their workloads are prioritized and gives the Cohesity Data Platform the information it needs to be able to efficiently run multiple, diverse, concurrent workloads (*Critical Capability #1*).

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Tying all of this functionality together is a modern, tiled UI. Here's a sample screenshot of the top level Cohesity Dashboard:

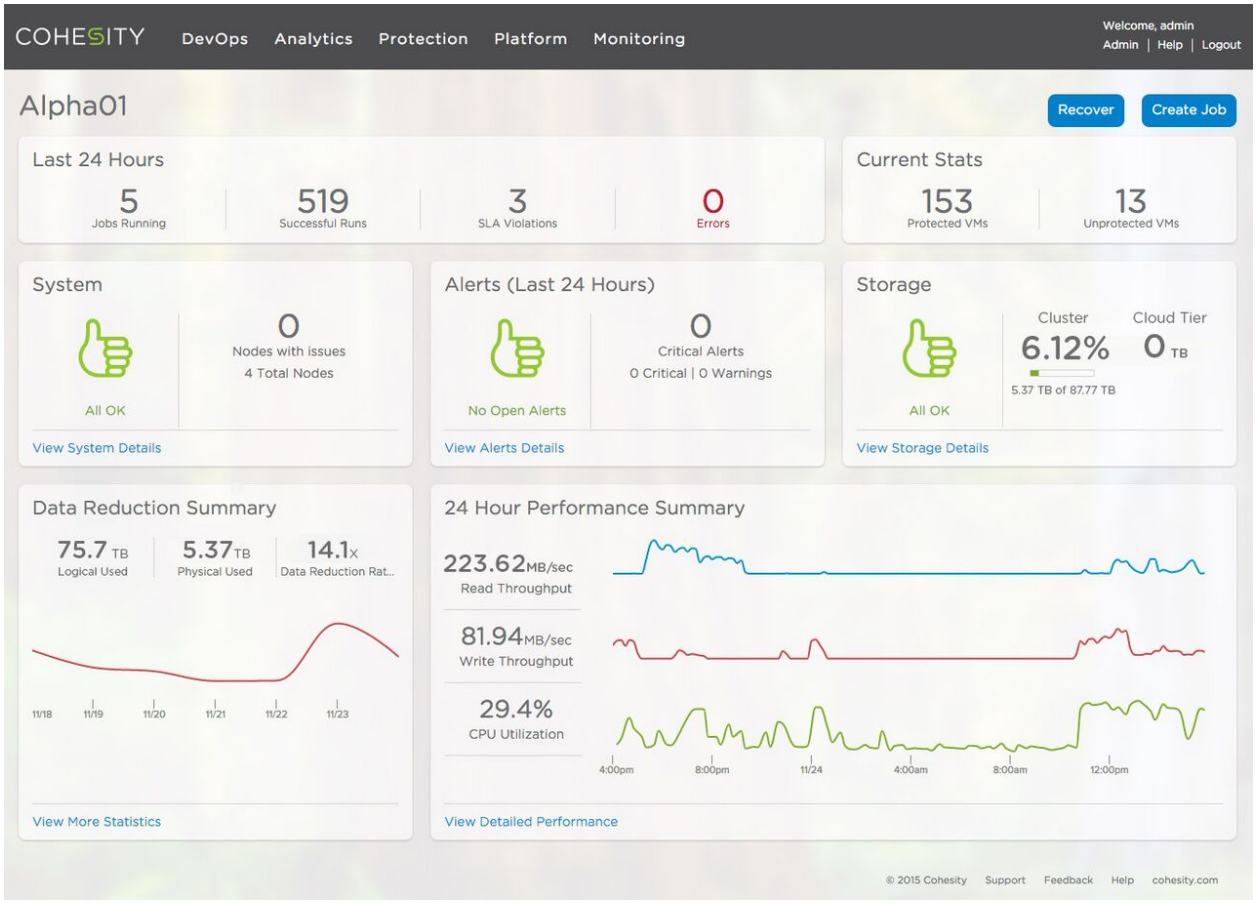

TANEJA GROUP OPINION

Traditionally, secondary storage was mostly a passive receptacle for copies of data that was originally homed in primary storage. The timing is now perfect to change this, because there's a confluence of technologies that weren't available ten years ago. Powerful, cost-effective, multi-core compute elements; efficient deduplication, compression and encryption algorithms; high capacity commodity HDDs; cost-effective Flash; highly evolved webscale file systems; virtualization applied to the entire IT stack; these are things that have already had dramatic impacts on the compute/primary storage space, but have yet to trickle down in an analogous way to secondary storage.

In truth, secondary storage has remained a mess. Yes, there have been many advances made over the years. Things such as disk-based backup storage, deduplication appliances, copy data management products and integrated backup appliances have all delivered solid improvements. But, these were individual attempts to solve particular standalone problems, not an effort to completely change the landscape of secondary storage as hyperconvergence has done in the compute/primary storage space. So, while the basic technology building blocks to resolve the fragmentation, inefficiencies and dark data issues of secondary storage now exist, what has been missing was the visionary thinking needed to pull them all together. Cohesity, as the pioneer in this new category of Hyperconverged Secondary Storage, is providing just that. Instead of having to manage multiple solutions from a variety of vendors, while trying to somehow stay on top of a rapidly growing mountain of data, customers can now select a simple, single vendor solution to their secondary storage problems – the Cohesity Data Platform.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Cohesity providing built-in data analytics out of the chute is goodness. Delivering them on top of the Cohesity Analytics Workbench, an open application integration platform with a Map Reduce framework and an SDK as part of their offering, is greatness. It will allow Cohesity, their customers and 3rd party developers to easily add powerful, custom-built analytics modules to their platform, significantly enhancing its flexibility and value. This is a key feature and significant differentiator for them. We expect it to be very popular.

Going forward, we believe that more hyperconverged secondary storage platform vendors are likely to appear; it's a big space, after all. It's almost certain that these platforms will provide even broader coverage of secondary storage workflows, such as adding data protection capabilities for Hyper-V, KVM and databases, which is a direction in which Cohesity is already heading. We also feel that Cohesity has set the category baseline by providing the hooks needed for anyone to build custom analytics on top of their platform; this is likely to become a "must have" feature in this space, so look for others to emulate it.

Cohesity is moving secondary storage away from its traditional passive role into a much more active space. This is only possible with deep understanding of how to build web scale architectures that can start small and grow almost infinitely; application of virtualization principles across the stack; building a file system that can deal with billions of files and a wide variety of unstructured and structured data; incorporation of inline compression and data deduplication; integration of self healing and high availability technologies; solid understanding of secondary storage use cases, and much, much more. It is not a challenge to be taken on by the faint hearted. We think the team at Cohesity is up to the task. Their architectural work is done but only the basic feature set is yet exposed. Based on what we have seen so far, the impact of this new architectural approach will potentially play havoc on traditional, legacy data protection vendors, as they struggle to deal with scale and lack of insight.

Cohesity's CDP is a paradigm shift, focused on simplification and flexibility while simultaneously shifting the cost curve down and enabling IT to efficiently squeeze maximum value out of all of business data assets. We're excited about their vision and we're optimistic about the path that they've taken.

Cohesity is not just hyperconverging secondary storage, they're redefining it.