## Point Solutions Have Resulted in Fragmentation, Overhead, and Added Complexity

The typical enterprise data center usually divides data storage into two primary categories: primary storage, which generally serves mission critical applications, and secondary storage, which houses workloads such as data protection, test and development, analytics, and general file services. To accomplish the various workloads included in secondary storage, businesses have adopted a variety of single-purpose solutions, each serving a particular business need, but resulting in unwanted data sprawl and increased complexity of management.

Today's data center administrators must constantly balance the needs of developers and the ever-growing volume of data they create. Every time a bug needs troubleshooting or an application needs a revision, oftentimes, the admin needs to make a copy of production data to either reproduce a problem or act as a real-world environment to code against. With the passive nature of traditional secondary storage solutions, admins generally must require a separate infrastructure stack to perform their testing and development, further increasing data sprawl across the enterprise.

Cohesity DevOps Prompts Data Consolidation with Ease

With Cohesity's intelligent secondary storage platform, virtualization administrators are able to leverage secondary storage workloads, such as data protection and management, to power test and development environments, drastically reducing their overall storage footprint. Cohesity's Open Architecture for Scalable, Intelligent Storage (OASIS) empowers developers to clone the latest backup of their production application stack and run it directly off the Cohesity Data Platform, providing a unified foundation for copy data management.
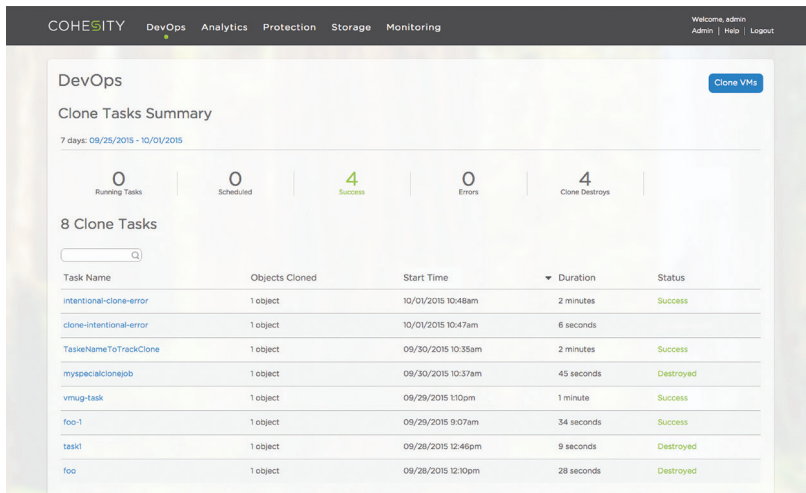
SnapTree™-Powered DevOps Environments Increase Business Agility and Time to Development

Cloning operations used for rapidly spinning up DevOps environments are powered by Cohesity's intelligent SnapTree™ data structures and snapshot capabilities. SnapTree™ is Cohesity'a patented tree-like data structure that enables a virtual, fully-hydrated image of every snapshot that is created. Since SnapTree™ clones are designed with performance in mind, access time is identical regardless of the number of snapshots or clones that are taken. This is radically different from legacy snapshot architectures,

Key Benefits:

- Reduce data center fragmentation by leveraging backups to drive DevOps workloads
- Enable self-service environments for developers and application owners to create instant test and development environments
- Instantaneously clone VMs as frequently as needed without performance penalties
- Transform idle backups into live data for DevOps

One Platform. Infinite Possibilities.

COHESITY

which rely on a chain-linked data structure. With SnapTree™, read operations don't have to traverse every link in the chain to find the requested piece of data, making snapshots and clones instantaneous.

In the case of test and development, an additional snapshot is taken when the clone is made, and presented back to VMware as a new unique virtual machine. All reads requested to old data are read directly from the older snapshots by following the short hops in the tree while all new or changed blocks are written to the new clone instance. This provides a full clone of live data with zero overhead that can be given to developers to perform their testing.

To facilitate rapid DevOps, developers and IT operations administrators have full access to the exposed REST API, which allows them to refresh the test and development environments automatically, either on a scheduled basis or on demand.

Cohesity's Flash-Accelerated Platform Significantly Reduces the Cost of DevOps Environments

The Cohesity C2000 Series provides the optimal combination of hard disk and PCI-e flash to deliver accelerated performance at the cost of secondary storage. As writes are sent to the Cohesity Data Platform, they are captured in a write cache and are eventually flushed to disk if the data gets cold. Cohesity's native intelligent tiering capabilities asses the individual profiles of each and every operation on the platform to strategically place data for quick access. If the I/O profile of the write is random in nature, the migration from cache to flash tier is instant. If the I/O profile is seen as sequential in nature, the flash tier of storage is skipped entirely and the data is sent straight to HDD. This ensures the PCI-e flash is used to accelerate the operations of data that is needed while other data is cost-effectively stored on hard disk. With Cohesity, DevOps engineers can benefit from the performance of semi-production storage at a fraction of the cost.

Reducing complexity while saving time and money is a priority for all CIOs. With Cohesity's comprehensive secondary storage solution, CIOs never have to choose between cost and performance. Cohesity enables IT departments to transform idle backups into live data for test and development.