Datrium™

Software Features Datasheet

At the heart of Datrium DVX is powerful DVX Software that simplifies management, scales the system from small to large with great performance, handles all cloud data management including scalable local, remote and cloud backup, and delivers the economics only available from Open Converged Infrastructure.

Split Provisioning

Inspired by post-2010 hyperscalers, Split Provisioning provides an approach to deliver either compute or capacity independently and incrementally for simpler administration and reduced costs.

Start with a single Data Node capable of supporting up to 100 effective terabytes of data and scale up to 10* Data Nodes and 1 PB of capacity. Start with a single Compute Node capable of an average of ~100,000 IOPS and scale to 128 Compute Nodes with more than 12 million total IOPS. Each Compute Node supports VMware vSphere, Linux KVM or Linux Containers.

Hosts are all stateless so any number of hosts or host SSDs can be down simultaneously, yet data is always safe and accessible with the DVX off-host Data Node. DVX Compute Nodes remain stateless and easily serviceable–much simpler for enterprise administrators versus Hyper-Converged.

Strong host IO isolation eliminates host-to-host neighbor noise, even when scaling and mixing diverse workloads.

Provision workload-specific hosts without neighboring host workloads affecting desired SLAs.

Scale up to 6,000 VMs or 60,000 objects where an object is a vDisk, file or container.
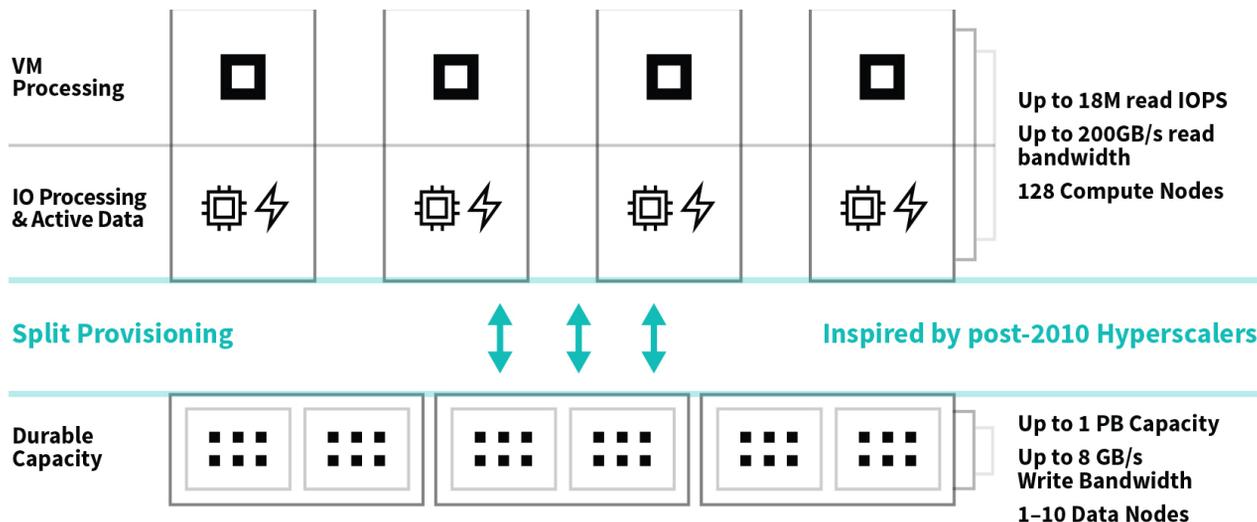
Insane Mode

Dynamically boost VM and Container storage performance by doubling the number of cores reserved for storage IO on a host. Quick and easy fix for workloads with heavy storage IO requirements.

Always-On Distributed Erasure Coding, Global Deduplication and Compression

All the storage efficiency and security benefits of erasure coding implemented in a distributed system. Used for all workloads all the time with no impact on performance.

Split Provisioning — figure 1

		Up to 18M read IOPS
VM Processing		Up to 200GB/s read bandwidth
IO Processing & Active Data		128 Compute Nodes

Split Provisioning — **Inspired by post-2010 Hyperscalers**

		Up to 1 PB Capacity
Durable Capacity		Up to 8 GB/s Write Bandwidth
		1–10 Data Nodes

Always-on erasure coding, compression, deduplication and other data services eliminate complexity and painstaking planning. No need to worry about hot vs. cold tiers, erasure coding vs. triple mirroring, compression vs. deduplication vs. no data reduction.

Ultra-efficient erasure coding, as well as compression and global deduplication across all hosts, provides the smallest durable data footprint possible (2x to 6x data reduction typical, often higher for VDI).

Commodity flash with no protection overhead (read-only), always-on data reduction for hot / cold data: Ultra-efficient erasure coding, inline deduplication and compression on flash per host, combined with global deduplication across all hosts.

Data Cloud Foundation

Data Cloud Foundation provides built-in, scalable and highly efficient copy data management. Near-instantaneous snapshots can be performed at datastore file -level granularity. Snapshots are completely self-contained and can be archived indefinitely and used for backup, archive, and DR. Utilize the Datrium VSS provider for application-consistent snaps that do not require any stunning of VM's.

Protection Groups with Dynamic Policy Binding enable users to create policy-based, automated groupings of snapshots.

Quickly and easily clone or restore snapshots, datastore files (virtual disks, VM templates, ISOs, OVAs, …), or Protection Groups.

Elastic Replication and Cloud Backup

Leverages processing, networking and local flash resources from all hosts within a cluster to replicate data concurrently, increasing the overall available replication bandwidth. Can tolerate multiple host failures and continue replicating data.

With Elastic Replication, you can maintain a hierarchy of local snapshots, remote snapshots on DVX and snapshots on AWS S3.

Blanket Encryption

Blanket Encryption is a 100% software-based solution that combines always-on efficient data reduction technology with high-speed encryption end-to-end, protecting data in-use at the host, in-flight across the network and at-rest on persistent storage.

Data is reduced, compressed and then encrypted as soon as it is created in the host RAM before it's written to the host flash or transmitted to the data server fully encrypted.

The internal key management service provides password & key rotation, lockdown mode that prompts for a passphrase upon a cold reboot before data can be accessed, and secure erase.

Adaptive Pathing

Smarter path aggregation for improved bandwidth utilization and resiliency.

DVX UI

Standalone HTML5 UI or vCenter HTML5 plug-in supports vSphere VMs, Linux VM's and Containers in a unified dashboard.

Provisioning capacity for a VM or Container is as simple as creating a vDisk from vCenter. No LUNs, no zones, no wondering what's making a VM's IO slow or how to fix it.

Simple, one-click software and firmware data node upgrades take typically 10 to 15 minutes and are done while your VMs remain entirely online.

With more than 6000 metrics designed into our system, you can track anything from IO latency of a single VM or SSD to Data Node fan speed to numerous other internal software metrics at any moment both in real-time and historically.

Integrated AutoSupport for 24x7 proactive monitoring and notifications.

Learn more at: https://www.datrium.com/datrium-dvx-software/