

Datrium DVX Datasheet

Datrium DVX Open Convergence Platform

DATRIUM DVX OPEN CONVERGENCE PLATFORM


The Datrium DVX platform converges storage and compute in a radical new way to enable a simpler journey to hybrid clouds, but without the rigidity of 3-tier convergence, and without the lock-in and scaling unpredictability of hyperconvergence (HCI). DVX separates on-host, software-driven IO services and performance from an off-host durable data repository, enabling N-host failure tolerance and host/host neighbor noise isolation for optimized performance in mixed-use clouds. DVX is open to almost any server type, including blades and 'brownfield' (existing compute infrastructure). Ultra-simple, VM-centric operation and analytics eliminate storage management. DVX makes what used to be the most expensive parts, now the low cost parts. The recipe for DVX affordability combines the use of existing and underutilized commodity server components with always-on, end-to-end dedupe and compression.

Seriously Fast.

- Performance scaling: Achieve higher performance and lower latency than all-flash arrays because reads never need to leave the host, avoiding SAN queuing delays. Adding a host adds up to 100,000 IOPS with sub-millisecond latency to DVX performance, versus dividing the shared resources of an array.
- Performance density: All read IO remains host-local in large Instance Flash, configurable up to 100 TB of effective flash/host, for low cost and ultra-low-latency IO operations.
- On-Demand Acceleration. IO processing leverages underutilized local CPU, so per-host performance can be doubled in-place and on-demand by reserving more available CPU resources (Insane Mode). If IO resources are strained on a host, simply vMotion to a host with more resource headroom to boost VM performance online.

Unbelievably Simple.

- VM-Centric: DVX ends storage management. Provisioning capacity for a VM 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.
- Automated Installation: Installation is 15 minutes or less from power-on to VM provisioning.
- One-Click Upgrades: Software upgrades are one-click, online and non-disruptive.



Incredibly Efficient.

- Open CPU and Flash: Leverage underutilized CPUs with installed server infrastructure, including both rack and blade servers. Take advantage of low cost, commodity server flash devices that are often one-eighth of the cost of array flash. The DVX platform includes licenses for its software on up to 32 hosts.
- Independent Scaling: Provision performance independent of capacity so costs stay low even with unpredictable growth or changing SLA needs.
- 2x – 6x effective data footprint reduction: Inline deduplication and compression on flash per host, combined with global deduplication across all hosts, provides the smallest data footprint possible (2x to 6x reduction typical, often higher for VDI).

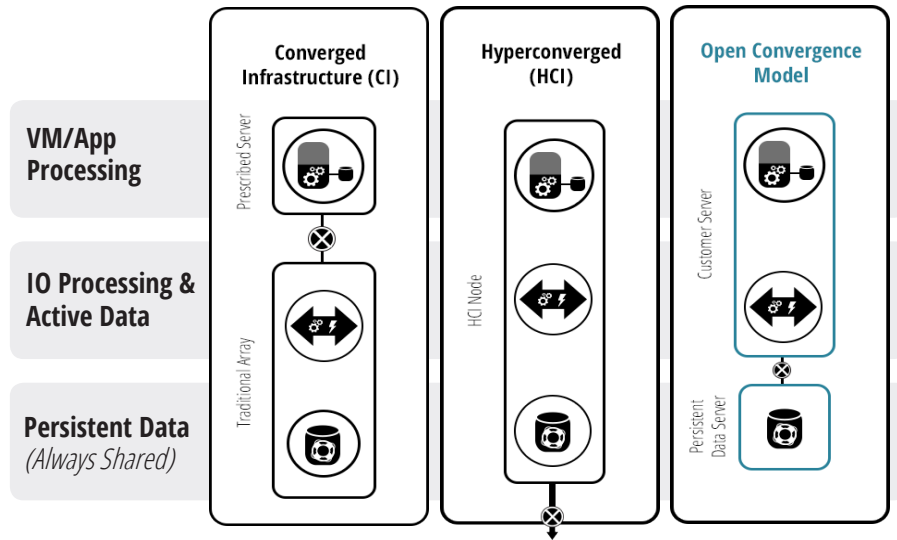
Powerfully Predictable.

- Mixed Workload Isolation: Strong host IO isolation eliminates host/host neighbor noise, even when mixing diverse workloads. Provision workload-specific hosts without neighboring host workloads affecting desired SLAs.
- Powerful VM Analytics: DVX IO and fault isolation vastly simplifies workload optimization and analytics as the environment scales, including an in-depth real-time view of all Host and VMs.
- N-Host Failure Tolerance: Any number of hosts can be down simultaneously, yet data is always safe and accessible with the DVX off-host data repository. DVX hosts remain stateless and easily serviceable--much simpler for enterprise administrators versus HCI or software defined storage. DVX includes 24x7 proactive monitoring and notifications.

OPEN CONVERGENCE FOR HYBRID CLOUDS

Datrium DVX solves the converged infrastructure problem in a revolutionary way, modeled on public cloud IaaS services instead of traditional 3-tier infrastructure or HCI. While recognizing that commodity servers and flash are faster and lower cost than arrays, DVX separates persistent data off-host in an HA data server. By aligning to an IaaS approach, DVX makes the journey to hybrid clouds much simpler over time.

DVX delivers higher IO performance and lower latencies than all flash arrays. It eliminates compromises with HCI, including server and configuration rigidity as well as noisy neighbor issues that have limited its fit for low-latency scaling and mixed workloads. DVX modularity simplifies granular provisioning and VM troubleshooting.



USE CASES

VDI / End User Computing

- Boot/login storms are handled in local flash, off-SAN, with ultra-low latency.
- Scale a few desktops to thousands with no/minimal latency increase.
- 4X faster desktop provisioning versus LUN based systems.
- Below \$2 per desktop per month when leveraging installed server infrastructure and below \$5 per desktop per month when deployed with new servers.



Private Cloud / VM Consolidation

- Scale performance and capacity independently
- Predictably lower latencies to VMs for reads on local flash versus all flash arrays.
- Configure hosts to specific workloads for IO performance, isolated from other hosts
- Grow VM IO speed on demand with Insane Mode or vMotion



Data Warehousing

- Flash-resident data enables up to 4x faster queries.
- Sequential query workload isolated from neighboring transactional applications for high performance for all DVX-based applications.



CUSTOMER QUOTES



I've been looking at the market for a long time, and this is really the first storage system that feels like it's built for an infrastructure-as-a-service provider like us.

–Martin Skojec, Director of IT



Northrim Bank

Datrium represents an evolutionary step in server-storage convergence by moving the logic, processing, and most of the touch points from the SAN Controllers to each Host Server, while maintaining worry-free data protection on their highly-available NetShelf.

–Benjamin Craig, EVP and CIO