



VX-CLOUD HYPER-SCALE ACCELERATED DATA PLATFORM

DATA SHEET

Delivering Tier 0 Performance at Petabyte Scale with the Lowest TCO

In order to extract real-time insights from massive data sets, infrastructure and operations teams architect systems for high velocity data access to run cognitive/AI systems as part of their digital transformation. In order to unleash the power of machine learning platforms, data must rapidly be ingested, inferenced and trained to achieve the stated business outcomes. Most business leaders evaluate Cloud infrastructure as a way to achieve the stated goals, but when low latency is a priority, public Cloud services are not always a cost-effective option. Most Cloud Service Providers (CSP's) do offer premium SSD based, performance tiers, unfortunately when data sets grow beyond 100TB, the cost of utilizing high performance Cloud service tiers to support these high-performance workloads can exceed \$1M annually. When faced with these options, these line of business leads will revert to their infrastructure teams to keep these services on-premises and build their own private cloud.

The challenge that the hyper-scale CSPs must overcome is how to deliver low-latency, petabyte and exabyte scale services at dramatically lower cost points to continue to offer cost-competitive services that deliver the right levels of performance and scalability. This is the opportunity that Vexata is targeting with the VX-Cloud Hyper-Scale Data Acceleration Platform.

VX-Cloud Hyper-Scale Data Acceleration Platform

VX-Cloud is a full software solution that utilizes industry standard servers and accelerated co-processors offering cloud scale block and file services for high performance compute (HPC), machine learning (ML) and Artificial Intelligence (AI) workloads. Complementary to the Vexata VX-100 Scalable NVMe Flash Array, VX-Cloud is purpose-built for managed service providers and enterprises looking to build cloud-scale services processing multi-petabyte data sets consisting of random, mixed read and write I/O. With VX-Cloud, large scale deployments of low-latency, high throughput services can be deployed as an Enterprise Private Cloud or deployed within a Cloud data center deployed as part of a high performance Infrastructure-as-a-Service (IaaS) offering.

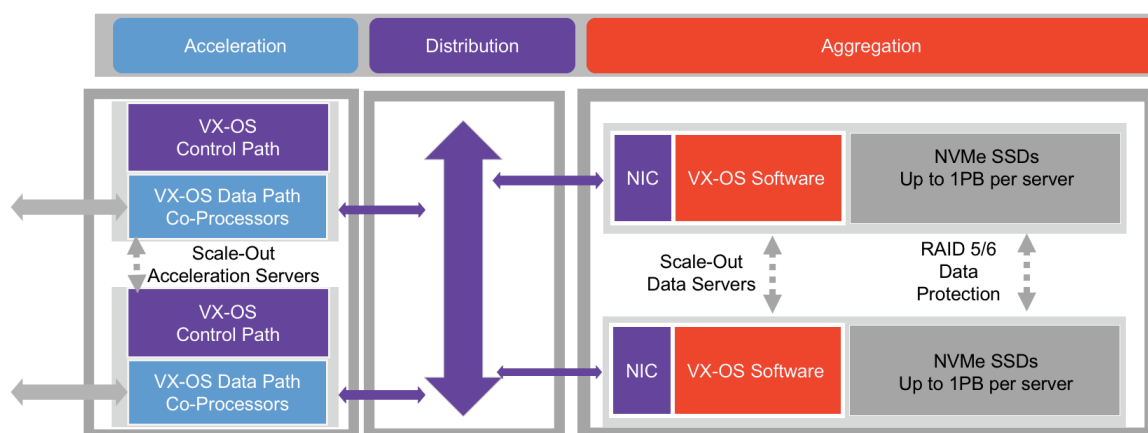


Figure 1: VX-Cloud Architecture

VX-Cloud consists of three core architectural elements, Acceleration, Distribution and Aggregation, each of which can scale independently on-demand. This is a key differentiation for VX-Cloud, which allows large enterprise and cloud providers to elastically scale performance and capacity independently, based on the requirements of the specific workload. VX-Cloud is delivered as a reference architecture through a network of strategic partners that deploy the software using best-in-class server platforms across both the aggregation and acceleration stages.

ABOUT VEXATA:

Founded on the premise that every business is challenged to deliver cognitive, data-intensive applications, Vexata delivers 10x performance AND efficiency improvements at a fraction of the cost of existing all-flash storage solutions. Learn more at www.vexata.com



VX-CLOUD HYPER-SCALE ACCELERATED DATA PLATFORM

DATA SHEET

VX-Cloud Accelerated Data Platform Specifications

Table 1. Software components

Software package	Associated Hardware	Features	Details
Data Acceleration Engine	Vx-DAE Standard Servers (Table2)	OS Support	Linux, Windows, ESX, Hyper-V
		Performance	1.25M IOPS and 10GB/s mixed R/W performance per license
		Scalability	10M IOPS and 80GB/s mixed R/W performance (current limit)
		Resiliency	High Availability, RAID5/RAID6, Multi-pathing, Non-disruptive Upgrades (NDU)
		Data Services	Thin provisioning, Space Efficient Snapshots and Clones, Data-at-Rest Encryption, Pattern Removal
Data Store Manager	Vx-DSM Standard Servers (Table2)	OS Support	Linux
		Performance	5GB/s mixed R/W performance per server
		Scalability	10's of Petabytes managed with price/TB licensing
		Data Services	Media Optimization, Garbage collection, Metadata management
Cloud Management Module	Vx-DAE Standard Servers	OS Support	Linux
		Interfaces	CLI, GUI, REST Api
		Management Services	Openstack, VMWare VSphere integration, Call home support, Cloud based monitoring

Table 2. Hardware Specifications (available through partner reference configurations)

Hardware	Features	Details
Data Acceleration Servers	Typical Server Config	TBD
	CPU sockets	2 or 4 sockets with Xeon Scalable Platform 18 core CPU per socket
	Memory	512GB (dual socket) 1024GB (quad socket)
	Accelerators	One per socket (Dual socket = 2, Quad socket = 4) Datacenter FPGA
	Host Interface	Up to 4 Adapters/NICS per server 2x40GE/2x100GE NVMe over Ethernet
	Storage Server (private) interface	2x100GE per Accelerator Up to 4 2x100GE interfaces per VxDAE server
	Data Store Manager	Typical Server Config
CPU Sockets		Dual socket with Xeon Scalable Platform 12 core CPU per socket
Memory		Minimum of 64GB DRAM + 128GB non-volatile memory (NV-DIMMs or similar)
DAE (private) interface		8x10GE/2x40GE/2x100GE Ethernet
Media		Up to 24 PCIe NVMe Flash SSDs per server 4TB, 8TB and 16TB SSDs supported
Switches	Lossless Layer2 Switch with 100GE support such as Mellanox SN3800	

Table 3. Performance & Capacity Scaling Specifications

DAE Server Type	Dual socket server						# DSM Servers in cluster	4	8	12	16
	2	2	3	3	4	4					
# DAE Servers in cluster	2	2	3	3	4	4	4 - 24	4 - 24	4 - 24	4 - 24	4 - 24
# (CPU, Accelerator) pairs per server	1	2	1	2	1	2	128TB - 768TB	256TB - 1536TB	512TB - 3172TB	1024TB - 6344TB	1024TB - 6344TB
Mixed 70/30 R/W IOPS	2.5M	5.0M	3.75M	7.5M	5M	10M	20GB/s	40GB/s	60GB/s	80GB/s	80GB/s
Bandwidth	20GB/s	40GB/s	30GB/s	60GB/s	40GB/s	80GB/s					
Latency	<200us reads and <100us writes										