

## ScaleIO

**ESSENTIALS**

- Software-defined block storage
- Performance and capacity on demand
- Multiple deployment options
- Infrastructure agnostic
- Elastic architecture
- Data protection across the cluster
- Meet the needs of traditional and next generation applications

**BENEFITS**

- Low TCO
- Scales to thousands of nodes
- Tens of millions of IOPS

Try it for free! Download the [Free & Frictionless version](#) today!

# EMC SCALEIO

## Software-Defined, Scale-Out SAN

### CHALLENGES

The management of large-scale, rapidly growing infrastructures is a constant challenge for many data center operation teams and it is not surprising that data storage is at the heart of these challenges. The traditional dedicated SAN and dedicated workloads cannot always provided the scale and flexibility needed. A storage array can't borrow capacity from another SAN if demand increases and can lead to data bottlenecks and a single point of failure. When delivering Infrastructure-as-a-Service (IaaS) or high performance applications, delays in response are simply not acceptable to customers or users.

### EMC SCALEIO: PUBLIC CLOUD AGILITY WITH PRIVATE CLOUD RESILIENCY

EMC® ScaleIO® is software that creates a server-based SAN from local application server storage to deliver flexible and scalable performance and capacity on demand. It converges storage and compute resources of commodity hardware into a single-layer architecture, aggregating capacity and performance, simplifying management, and scaling to thousands of nodes. As an alternative to a traditional SAN infrastructure, ScaleIO combines HDDs, SSDs, and PCIe flash cards to create a virtual pool of block storage with varying performance tiers. In addition, it provides enterprise-grade data protection, multi-tenant capabilities, and add-on enterprise features such as QoS, thin provisioning, and snapshots. ScaleIO is hardware-agnostic, supports physical and/or virtual application servers, and has been proven to deliver significant TCO savings vs. traditional SAN.

### MASSIVE SCALABILITY

ScaleIO is designed to massively scale from three to thousands of nodes. Unlike most traditional storage systems, as the number of storage devices grows, so do throughput and IOPS. The scalability of performance is linear with regard to the growth of the deployment. Whenever the need arises, additional storage and compute resources (i.e., additional servers and/or drives) can be added modularly so that resources can grow individually or together to maintain balance. Storage growth is therefore always automatically aligned with application needs.

### EXTREME PERFORMANCE

Every server in the ScaleIO cluster is used in the processing of I/O operations, making all I/O and throughput accessible to any application within the cluster. Such massive I/O parallelism eliminates bottlenecks. Throughput and IOPS scale in direct proportion to the number of servers and local storage devices added to the system, improving cost/performance rates with growth. Performance optimization is automatic; whenever rebuilds and rebalances are needed, they occur in the background with minimal or no impact to applications and users. The ScaleIO system autonomously manages performance hot spots and data layout. EMC lab testing results demonstrate:

## Features Essential for Enterprises and Service Providers

- Partitioning / Protection Domains
- Tiering
- Multi-tenancy
- Snapshots
- Fault set domains
- Two-Way Mesh Mirroring
- Storage QoS
- Integrations with VMware™ vSphere® and OpenStack®

## NEW ESSENTIAL FEATURES IN SCALEIO 2.0

- In-flight Checksum
- 5 Node MDM Cluster
- Read Flash Cache
- EMC Secure Remote Services (ESRS)
- Instant Maintenance Mode
- Software Component Validation
- CoreOS Support
- Ubuntu Support
- OpenStack Support

Workload	IOPS (3 Nodes)	IOPS (128 Nodes)
<b>100% READ</b>	~875,000	~31,000,000
<b>70% READ / 30% WRITE</b>	~650,000	~23,750,000
<b>100% WRITE</b>	~375,000	~12,500,000

## UNPARALLELED FLEXIBILITY

ScaleIO provides flexible deployment options. With ScaleIO, you are provided with two deployment options. The first option is called “two-layer” storage-only and is when the application and storage are installed in separate servers in the ScaleIO cluster. This provides efficient parallelism and no single points of failure. The second option is called “hyper-converged” and is when the application and storage are installed on the same servers in the ScaleIO cluster. This creates a single-layer architecture and provides the lowest footprint and cost profile.



ScaleIO provides unmatched choice for these deployment options. ScaleIO is infrastructure agnostic making it a true software-defined storage product. It can be used with mixed server brands, operating systems (physical and virtual), and storage media types (HDDs, SSDs, and PCIe flash cards). In addition, customers can also use OpenStack commodity hardware for storage and compute nodes.

## SUPREME ELASTICITY

With ScaleIO, storage and compute resources can be increased or decreased whenever the need arises. The system automatically rebalances data “on the fly” with no downtime. Additions and removals can be done in small or large increments. No capacity planning or complex reconfiguration due to interoperability constraints is required, which reduces complexity and cost. The ScaleIO system reconfigures itself as the underlying resources change; data is rearranged and spread evenly on the servers to optimize performance and enhance resilience. All of this happens automatically without operator intervention and therefore eliminates the need for costly and disruptive data migrations.

## ESSENTIAL FEATURES FOR ENTERPRISES AND SERVICE PROVIDERS

ScaleIO offers a set of features that gives you complete control over performance, capacity and data location. For both private cloud data centers and service providers, these features enhance system control and manageability—ensuring that quality of service (QoS) is met. With ScaleIO, you can limit the amount of performance—IOPS or bandwidth—that selected customers can consume. The limiter allows for resource distribution to be imposed and regulated, preventing application “hogging” scenarios. Data masking can be used to provide added security for sensitive customer data. ScaleIO offers instantaneous, writeable snapshots for data backups.

“...we've never seen any HA shared storage perform as fast as ScaleIO has in our lab. ScaleIO has shattered records in every situation so far, pushing our testing infrastructure to its limits.”

StorageReview Enterprise Lab - “EMC VxRack Node powered by ScaleIO Review”, Feb 2016

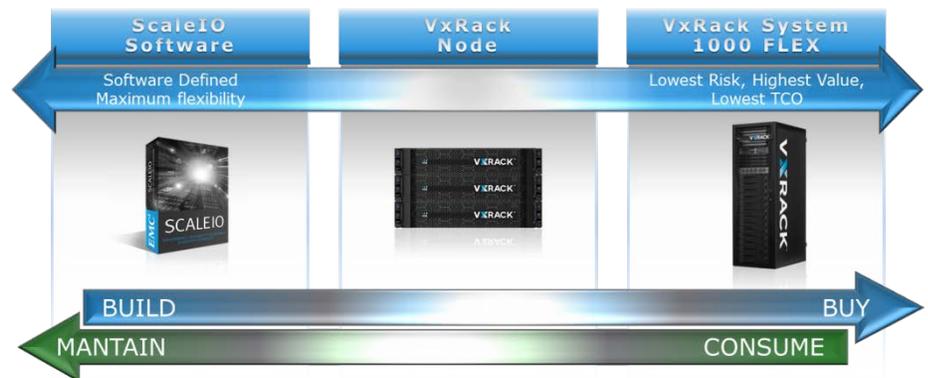
For improved read performance, DRAM caching enables you to improve read access by using SDS server RAM. Fault sets – a group of SDS that are likely to go down together – can be defined to ensure data mirroring occurs outside the group, improving business continuity. You can create volumes with thin provisioning, providing on-demand storage as well as faster setup and startup times.

ScaleIO also provides multi-tenant capabilities via protection domains and storage pools. Protection domains allow you to isolate specific servers and data sets. This can be done at the granularity of a single customer so that each customer can be under a different SLA. Storage pools can be used for further data segregation, tiering, and performance management. For example, data that is accessed very frequently can be stored in a flash-only storage pool for the lowest latency, while less frequently accessed data can be stored in a low-cost, high-capacity pool of spinning disks.

## CHOOSE THE RIGHT SOLUTION FOR YOU

When it comes to SDS, whether prefer a “build it yourself” or “turnkey” approach to your solution, ScaleIO delivers the consumptions that will help you reach your goals quickly and efficiently. You have the option to consume ScaleIO in 3 different models:

- ScaleIO software – delivers the maximum flexibility and installs directly on your x86 standard server hardware and infrastructure.
- VxRack Node – bundles x86 commodity servers and ScaleIO software, enabling you to quickly deploy a fully architected software-defined, scale out, server SAN
- VxRack System 1000 FLEX - a fully integrated, turnkey, hyper-converged solution which leverages ScaleIO, enabling you to hit the ground running.



## CONTACT US

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, [contact](#) your local representative or authorized reseller, visit [www.emc.com](http://www.emc.com), or explore and compare products in the [EMC Store](#).

EMC2, EMC, the EMC logo, ScaleIO, are registered trademarks or trademarks of EMC Corporation in the United States and other countries. © Copyright 2016, 2015 EMC Corporation. All rights reserved. Published in the USA. 07/15, Datasheet H12713.5.

VMware and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

EMC<sup>2</sup>