

# Modernizing Virtual Infrastructures Using VxRack FLEX with ScaleIO

Date: September 2017 Author: Mike Leone, Senior Validation Analyst | Dom Amato, Associate Validation Analyst

## Background

As organizations continue to look for ways to modernize their infrastructures by delivering a cloud-like experience on-premises, hyperconverged offerings are exceeding expectations. In fact, the adoption of hyperconverged infrastructure has more than doubled over the last year and shows no signs of slowing down.<sup>1</sup> Simplified, flexible deployment options that are easily managed and easily scale are just a few of the factors driving the rapid adoption of hyperconverged technology. As consolidation and modernization efforts are well underway across all of IT, hyperconverged offerings continue to meet the needs of the business while satisfying existing and future application SLAs, whether they are related to performance, scalability, reliability, or cost.

With many organizations running enterprise applications and databases of different shapes and sizes, including hypervisors and operating systems, on traditional physical or virtual infrastructure models, all configurations and architectures must be considered when planning and expanding further adoption of hyperconverged infrastructure in the data center.

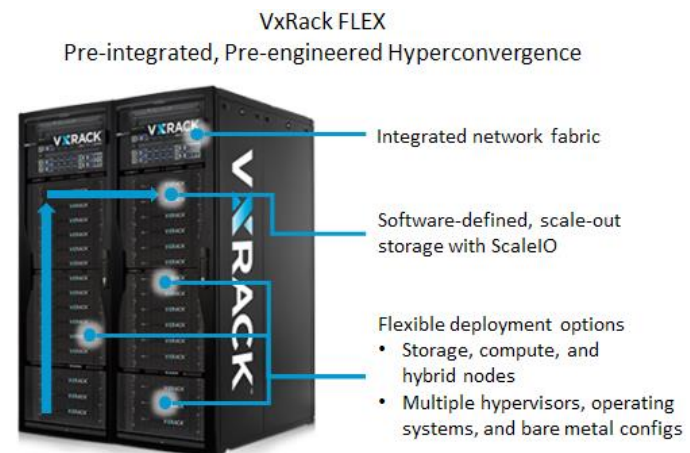
## VxRack System FLEX

VxRack System FLEX (VxRack FLEX) is a rack-based hyperconverged solution that leverages Dell EMC ScaleIO to deliver flexible, scalable, and performant software-defined storage. The ScaleIO software runs on industry-proven Dell EMC PowerEdge servers to deliver a full hyperconverged stack that enables IT organizations to deal with one vendor that can supply all the software, hardware, and support required to modernize their data centers. Key to the solution is the scale-out architecture, offering organizations flexibility to start small and grow based on their needs. Further, this elasticity delivers on the hyperconverged promise of cloud-like scale and flexibility on-premises.

While many other hyperconverged offerings disregard networking altogether, VxRack FLEX supports both physical and virtual networking and includes top of rack switches that control network traffic, management, and redundancy. Put it all together and organizations get a hyperconverged solution that can easily be used with all other Dell EMC products and services, together with ScaleIO delivering impressive performance and protection at scale. In addition to ScaleIO data protection capabilities, advanced data protection, replication, availability, and recovery are available with technologies such as Data Domain with Data Domain Boost and RecoverPoint.

## ScaleIO – Data Center Grade Software-defined Storage

Powering the VxRack FLEX solution is ScaleIO, Dell EMC's scale-out software-defined storage solution that abstracts the direct-attached storage found in Dell EMC PowerEdge servers into a pool of shared block storage. By converging the storage and compute on the same physical servers, this single and/or two-layer architecture helps to simplify management and maximize storage efficiency as the infrastructure grows from four to thousands of nodes. Whether using HDDs, SSDs, or even NVMe or PCIe flash, storage is combined into virtual block-storage pools with varying performance tiers. Combined with QoS, snapshots, caching, fault sets and protection domains, and data-at-rest encryption, ScaleIO running within the VxRack FLEX system delivers a data-center-grade, fully-integrated, hyperconverged solution. Leveraging a software-defined storage approach to satisfy enterprise application and database block storage requirements enables organizations to potentially break free of large initial investments and high operational costs commonly associated with traditional SANs. Further, fears of technology updates, refreshes, and data migrations impacting costs, risk, and periods of downtime can be all but eliminated.



<sup>1</sup> Source: ESG Research Report, *Hyperconverged Infrastructure Continues to Gain On-premises Momentum*, to be published.

This ESG Lab Review was commissioned by Dell EMC and is distributed under license from ESG.

## Traditional Virtual Infrastructures with SANs


Virtualization has had a profound impact on enterprise IT by enabling infrastructure consolidation and, therefore, cost reductions. With converged infrastructures, virtualization allows for a number of improved efficiencies by providing a pre-configured, pre-tested package of compute, networking, and storage. However, these traditional infrastructures use SAN-driven storage, which comes with numerous limitations in a virtual environment. Not only is SAN provisioning time-consuming, but its configuration flexibility is limited. Allocating workloads within a SAN is complex and supplementing with additional arrays to ensure predictability is costly and far from a guarantee of success. Further, converged infrastructures may add costly delays during upgrades while also limiting customers' choices when it comes to vendors. This is true not just for the physical components of the converged infrastructure but for the virtualization layer.

Considering the limitations of infrastructures that depend on a traditional SAN, software-defined storage (SDS) can be a more attractive option for organizations looking to modernize their infrastructures. SDS brings flexible, scalable storage to infrastructures, most recently as a component of hyperconverged implementations running a particular vendor's hypervisor, whether VMware, Microsoft Hyper-V, or KVM. In some cases, hyperconverged offerings may impose limits by not allowing storage and compute to be decoupled. Not only does this limit scaling options, but it also limits deployment options by forcing organizations to select one hypervisor. Organizations are looking for a flexible HCI offering with proven SDS capabilities that can support an organization's need for a heterogeneous hypervisor infrastructure.

## Modernizing Virtual Infrastructures with ScaleIO and VxRack FLEX

Dell EMC's VxRack FLEX offers a way to optimize traditional virtualized infrastructures by supporting multiple hypervisors, operating systems, and bare metal configurations. With ScaleIO, FLEX allows for independent scaling of compute and storage, eliminating stranded resources and bringing SDS flexibility to virtualized infrastructures that traditional SAN can not provide. Since ScaleIO software is inherently virtualized, its resource provisioning and adaptability to different hardware and software enable higher levels of flexibility. Scaling now becomes significantly easier with decoupled compute and storage, enabling an optimal distribution of resources based on application and workload needs.

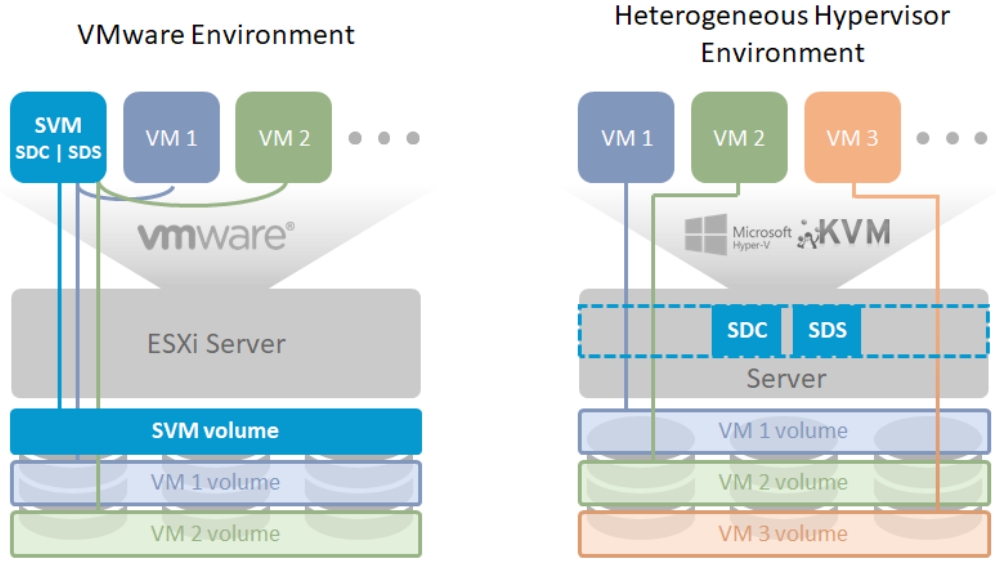
The figure demonstrates how VxRack FLEX allows for improved deployment flexibility and simpler resource provisioning with decoupled compute and storage in a virtualized infrastructure that leverages heterogeneous hypervisors. Several ScaleIO volumes are spread across storage-only nodes, with each volume associated with a hypervisor—VMware ESXi, Microsoft Hyper-V, and Linux KVM. With ScaleIO providing scale-out storage, regardless of where the storage sits physically, dynamic elasticity combats the typical issues with attempting to manage unpredictable, evolving workloads. Furthermore, should additional compute or storage be required, they can scale individually as needed to avoid excessive spend on unnecessary resources. ScaleIO is also resilient in the event of server or disk failure and can recover I/O operations without manual interference to minimize outages. VxRack FLEX with ScaleIO relieves the burden of capacity planning and recovery for virtualized data centers that are traditionally synonymous with expensive and error-prone maintenance.



## Virtual Architectures with ScaleIO and VxRack FLEX

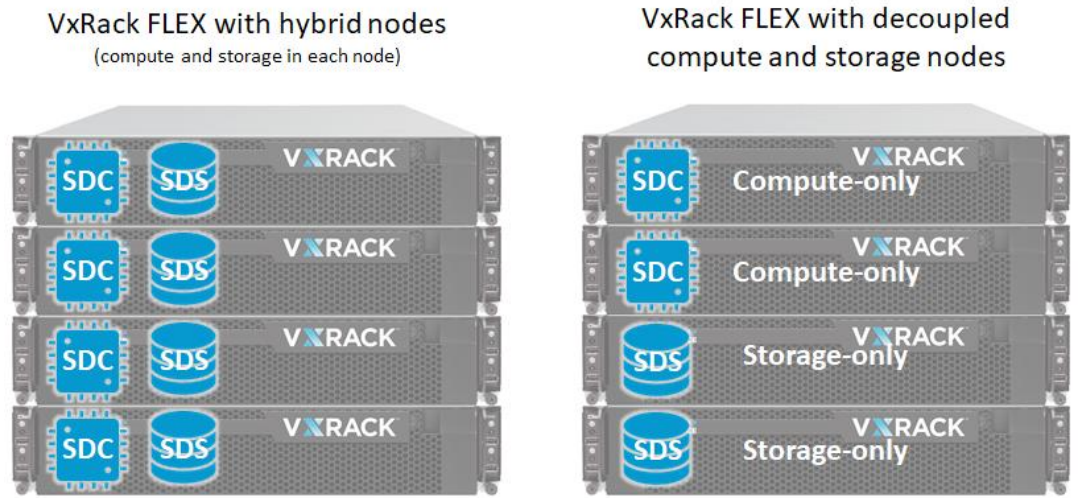
At the center of VxRack FLEX’s hyperconverged offering is ScaleIO, which consists of three software components that form the core of the architecture: ScaleIO Data Servers (SDS), ScaleIO Data Clients (SDC), and Meta Data Managers (MDM). The SDS manages local storage on a server that is part of the overall ScaleIO cluster and accesses any local storage space to fulfill I/O requests. The SDC distributes the I/O requests to the SDS and allows hypervisors and operating systems to access storage volumes within the cluster. The SDC is intelligent in that it knows which SDS to contact depending on where a request originates, with each SDC providing multi-path access across every SDS instance in the cluster. Lastly, MDMs serve as the control mechanism that maps clients to their data and automatically repairs operations following hardware failures. Together, these components run natively on ESX, Windows, or Linux to simplify deployments regardless of the running hypervisor.

In a traditional two-layer virtual infrastructure deployment with ScaleIO, where application servers and storage servers are separate, SDSs and SDCs run separately. SDCs run on the application servers, while SDSs run on the storage servers. MDMs are further removed from the architecture, only becoming active in the event of node failures or layout changes. As far as the hypervisor goes, ScaleIO can be deployed to support both VMware or non-VMware environments, such



as Microsoft Hyper-V and Linux KVM. As shown on the left side of the figure, in a VMware environment, ScaleIO provisions a dedicated VM that runs on each ESX host called the ScaleIO VM (SVM), which contains both the SDS and SDC. Each VM within the ESX server has access to each other through the SVM, while VMs within other ESX servers have cross-server access through the SDS. On the right side of the figure, non-VMware environments (Hyper-V and KVM) deploy as a component of the hypervisor itself, with both SDCs and SDSs installed on the physical server as opposed to within a VM.

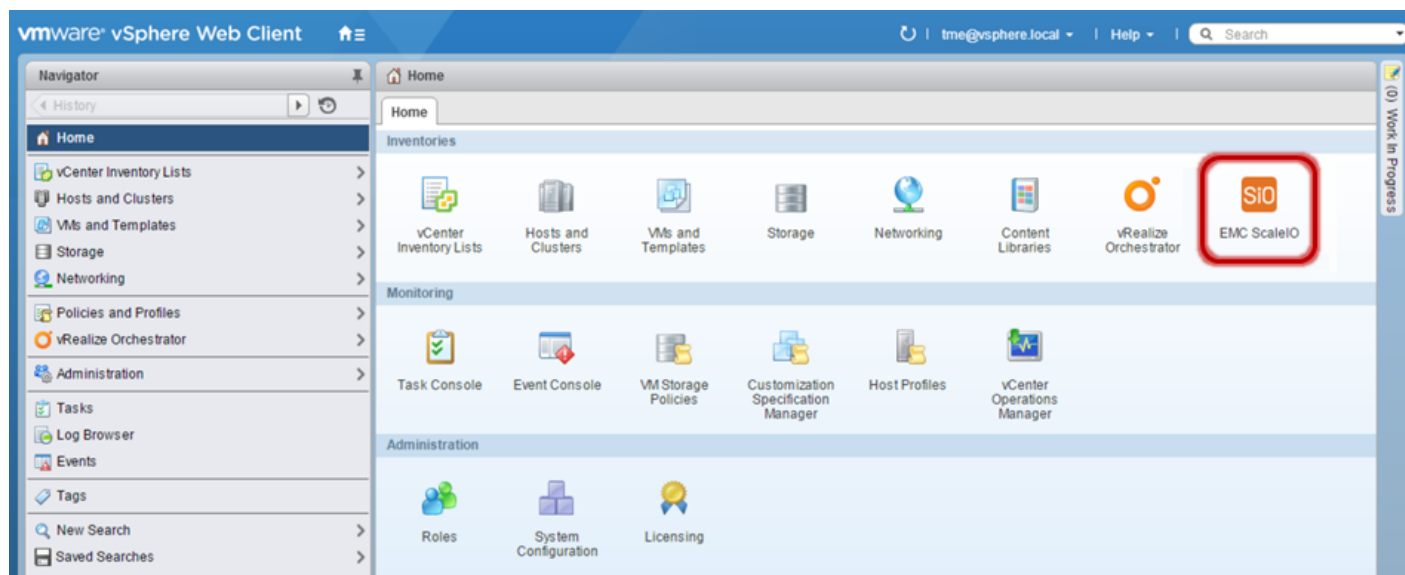
In VxRack FLEX’s hyperconverged infrastructure, the SDSs and SDCs are deployed differently depending on the type of node. For hybrid nodes that contain both compute and storage, the SDS and SDC reside on the same node, which optimizes resource utilization within the capabilities of the hardware. In these types of environments, where storage and compute are managed as one, organizations are able to yield higher levels of flexibility and performance. In the case of deploying compute- and storage-only nodes in the VxRack FLEX, SDSs are installed on the storage-only nodes, while the SDCs reside on the compute-only nodes.



## Installing ScaleIO with VMware

While VxRack FLEX with ScaleIO provides heterogeneous hypervisor support, ESG focused on VMware, walking through the deployment process using the VMware deployment plugin, which not only streamlines the installation of the software, but allows for more flexible resource provisioning. Users do have the option to manually deploy ScaleIO in ESX environments too.

First, the ScaleIO plugin must be registered with vCenter, which can be accomplished by extracting an OVA from the ScaleIO image. Once the appropriate vCenter identifiers are specified, “Standard” or “Advanced” registration options are available. ESG validated the standard procedure. Once a user has logged into vCenter and has downloaded and registered the ScaleIO plugin, it should be viewable in the vSphere interface. It should be noted that a best practice is to use the OVA template to create SVM templates within ScaleIO.



Next, SDCs must be installed on all ESX hosts. This process is very simple, as the “Basic” tasks section of ScaleIO shows the option to install. After clicking this option, simply select the desired hosts and click “Install” and “Finish” when completed. The last step in this phase of the installation is to restart each ESX host.

After installing the SDCs, ScaleIO can be deployed. VMware vCenter does allow for some advanced configuration, such as the option to decide between virtual machine disk (VMD) and raw device mapping (RDM) for storage formatting. A best practice is to increase the parallelism limit when deploying for larger-scale systems (greater than 100 nodes), enabling faster installation depending on available processing power. Selecting the option to deploy ScaleIO enables all ESX hosts to be added to the cluster, with a minimum requirement of three nodes. Additional configuration operations are available that are associated with three or five node clusters, individual component performance profiles, protection domains, and storage pools.

It should be noted that to enable the use of RecoverPoint, “zero padding” must be enabled for the storage pools. In addition, each ESX host contributing to Scale IO requires a protection domain and that the option to optimize for flash be turned on, if necessary. Finally, ESX hosts are mapped and assigned to their desired SDS component.

## The Bigger Truth

As organizations look for solutions to help modernize their infrastructure, hyperconverged infrastructures are serving as go-to architectures due to their underlying ability to improve time to value. With VxRack FLEX, customers receive a pre-integrated, pretested, and pre-validated hyperconverged infrastructure that can be operational within hours of arriving on the loading dock. With software-defined storage from ScaleIO, organizations gain a flexible underlying storage solution that enables linear performance scalability, while the VxRack FLEX node configuration flexibility enables organizations to easily scale out just compute, just storage, or both.

In order to take full advantage of virtualization technology, organizations need a flexible, software-defined, virtual infrastructure capable of surpassing the limitations of traditional virtual infrastructures. VxRack FLEX with ScaleIO is a modern hyperconverged solution capable of heterogeneous hypervisor support. Decoupled compute and storage not only enables flexible deployment options, but provides the scalability organizations demand while optimizing resource provisioning and utilization. This level of agility makes the VxRack FLEX an ideal candidate for organizations exploring modernization initiatives that must incorporate hyperconverged infrastructures.

ESG suggests exploring the VxRack FLEX offering as a way for organizations to future-proof their virtual IT infrastructures while not only continuing to meet the strict performance and protection requirements of their traditional mission-critical database applications, but also by delivering IaaS and meeting the requirements of emerging data analytics platforms.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.