

# OBJECTSCALE: NEXT EVOLUTION IN OBJECT STORAGE



## Nithin Shankar

Inside Product Specialist 2  
Dell Technologies  
N.shankar@dell.com

## Kshitij Yadav

Inside Product Specialist 2  
Dell Technologies  
Kshitij.yadav@dell.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

[Learn more at www.dell.com/certification](http://www.dell.com/certification)

# Table of Contents

1. INTRODUCTION .....	4
1.1 IDC .....	4
1.2 BACKGROUND .....	4
1.3 What is Object Storage? .....	4
1.4 Software-Defined Storage.....	5
2. ObjectScale.....	6
2.1 INTRODUCTION.....	6
2.2 ECS .....	6
2.2.1 ECS Platform.....	7
2.2.2 ECS Storage Structure.....	9
2.3 ObjectScale Architecture .....	11
2.3.1 Object Store.....	12
3. Deployment Options with ObjectScale .....	14
3.1 ObjectScale on Red Hat OpenShift .....	14
3.1.2 Architecture .....	14
3.1.3 Dell Technologies Services for ObjectScale on Red Hat OpenShift.....	15
3.2. Dell EMC ObjectScale on VMware vSphere with VMware Tanzu .....	15
3.2.1 Architecture .....	16
3.3 Dell EMC ObjectScale on VMware vSphere with Tanzu on Dell EMC VxRail .....	16
3.3.1 Advantages of ObjectScale on VxRail .....	16
4.ObjectScale Use cases.....	18
5. ObjectScale and ECS Positioning .....	20
6. CONCLUSION .....	21
7. REFERENCES .....	23

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

# 1. Introduction

## 1.1 IDC

According to IDC, by 2026, large companies forecast that unstructured data stored as files or object storage in on-premises, edge, or public clouds will triple compared to 2021. By 2025, 60% of the global storage capacity for unstructured data will be deployed as software-defined storage (SDS), up from less than 25% in 2021.

## 1.2 Background

The unstructured data storage market is evolving rapidly. New capabilities are needed to meet the growing challenges of exponential data growth, rapid digitization, and business globalization. Storage infrastructure products based on distributed file systems and object storage are growing faster than ever, both in scale and capacity, as a platform for meeting the growth of unstructured data in global data centers.

Gartner end users report that unstructured data has increased by more than 30% year-on-year. As a result, Infrastructure and Operations (I&O) leaders are looking for scalable on-premises storage products that can meet growing digital business use cases while reducing acquisition, operational, and management costs. I&O leaders want scalability, flexibility, lifecycle management, manageability, and analytical insights into the data.

Rapid growth of unstructured data for new and established workloads requires new types of products and cost efficiency. Most products in this market are driven by SDS that can provide petabytes of storage. SDS can also leverage hybrid cloud workflows with public cloud infrastructure as a service (IaaS) to reduce total cost of ownership (TCO) and improve data mobility. Newly established storage vendors continue to develop scalable storage scalable file systems and object storage products to address the cost, agility, and scalability limitations of traditional scale-up storage environments.

## 1.3 What is Object Storage?

Gartner defines file systems and object storage as a distributed design-based software and hardware platform that supports object and/or scale-out file system technology to meet the growing demands of unstructured data. This market is based on a distributed computing architecture with no single points of failure or contention throughout the system. More specifically, the product should have a fully distributed architecture. In this architecture, data and metadata are distributed, replicated, or erasure-coded across multiple nodes in the cluster. When managing a multi-petabyte scale system, it is important to be able to scale out the capacity and throughput of the cluster by adding independent nodes to the global namespace / file system.

Object storage, according to Gartner, also refers to systems and software that store data in "objects" and offer it to clients using RESTful HTTP (application programming interface) APIs, such as Amazon Simple Storage Service (S3), which has become the de facto standard for

object storage access.

## 1.4 Software-Defined Storage

SDS is a storage architecture that separates storage and software from the underlying hardware infrastructure. Unlike traditional network attached storage (NAS) or storage area network (SAN) systems, SDS is typically designed to run on industry standard systems, and the software does not rely on proprietary hardware.

Decoupling storage software from its hardware allows you to expand your storage capacity as needed, rather than adding another piece of proprietary hardware. You can also upgrade or downgrade hardware at any time. Basically, SDS offers a great deal of flexibility.

In most cases, SDS should have:

- **Automation:** Simplified management to keeps costs down.
- **Standard interface:** An API for management and maintenance of storage devices and services.
- **A virtualized data path:** Interfaces for block, file, and object that support applications written to these interfaces.
- **Scalability:** The ability to scale out storage infrastructure without impeding performance.
- **Transparency:** The ability to monitor and manage storage use while keeping track of available resources and costs.

Benefits of Software-defined Storage:

- **Flexibility in hardware selection:** The SDS you choose does not have to be from the same company that sold you the hardware. You can use any commodity server to build your SDS-based storage infrastructure. This means that you can maximize the capacity of your existing hardware as your storage needs grow.
- **Cost efficiency:** SDS is distributed and scales out instead of scaling up, allowing you to adjust capacity and performance independently.
- **You can join many data sources to build your storage infrastructure:** You can network object platforms, external disk systems, disk or flash resources, virtual servers, and cloud-based resources (even data dedicated to workloads) to create a unified storage volume.
- **SDS can adjust automatically based on your capacity needs:** SDS doesn't depend upon hardware; automation in SDS is, well, automatic in the sense that it can pull from any storage volume it's connected to. Without administrator involvement, new connections, or new hardware, the storage system can adapt to data needs and performance.
- **Limitless Scalability:** Traditional SAN are limited to the number of nodes (devices with assigned IP addresses) they can use. SDS, by its very definition, is not similarly constrained. That means – theoretically – it is infinitely scalable.

## 2. ObjectScale

### 2.1 Introduction

As organizations adopt containers and Kubernetes to run next-generation applications in the enterprise data center, they must also deliver a modern storage infrastructure that aligns the needs of developers and IT.

This is accomplished through use of a cloud-native operating architecture that allows infrastructure to be disaggregated so that computing and storage resources can be scaled separately to meet specific workload requirements. This capability is crucial as tremendous data expansion increases the need on IT to create storage designs that can support data-intensive artificial intelligence (AI), Internet of Things (IoT), and analytics apps, as well as traditional workloads at scale.

IT must provide developers with agile infrastructure that gives them more efficient ways to provision storage for their apps while ensuring the entire environment is secure and governed according to policy as containers and Kubernetes continue to drive cloud-native modernization initiatives. This necessitates S3-compatible, Kubernetes-compliant object storage that works in tandem with CI/CD (continuous integration and continuous deployment) pipelines to speed up application development and reduce time-to-market.

### 2.2 ECS

Dell Technologies has a long history in the field of object storage. In 2001, Dell developed Dell EMC Centera, the first CAS-based commercial object storage solution, and since then continually redesigned the portfolio to match customers' needs.

Elastic Cloud Storage (ECS), Dell EMC's flagship enterprise-class object storage platform, is designed to handle both traditional and next-generation workloads. ECS has unrivalled scalability, performance, robustness, and economics to meet the demands of a modern organization, and it can be deployed as a software-only architecture or as a turnkey appliance. With ECS, you can create a cost-effective object storage private cloud with TCO as low as 59 percent. As an S3-compatible, globally scalable object-store, ECS combines the capabilities of the public cloud with the command and control of a private cloud infrastructure.

ECS is a software-defined cloud storage platform that allows you to store, manipulate, and analyze unstructured data on commodity hardware at a large scale. It blends commodity infrastructure's low cost with traditional arrays' enterprise-level reliability, availability, and serviceability.

ECS has a scalable design with several nodes and storage devices attached. The nodes and storage devices are commodity components that are housed in one or more racks and are similar to devices that are widely accessible.

## 2.2.1 ECS Platform

In ECS, each function is implemented as a separate layer, resulting in a layered architecture. This design strategy ensures high availability by making each layer horizontally scalable across all nodes in the system. The ECS platform is made up of the data services, portal, storage engine, fabric, infrastructure, and hardware component layers.




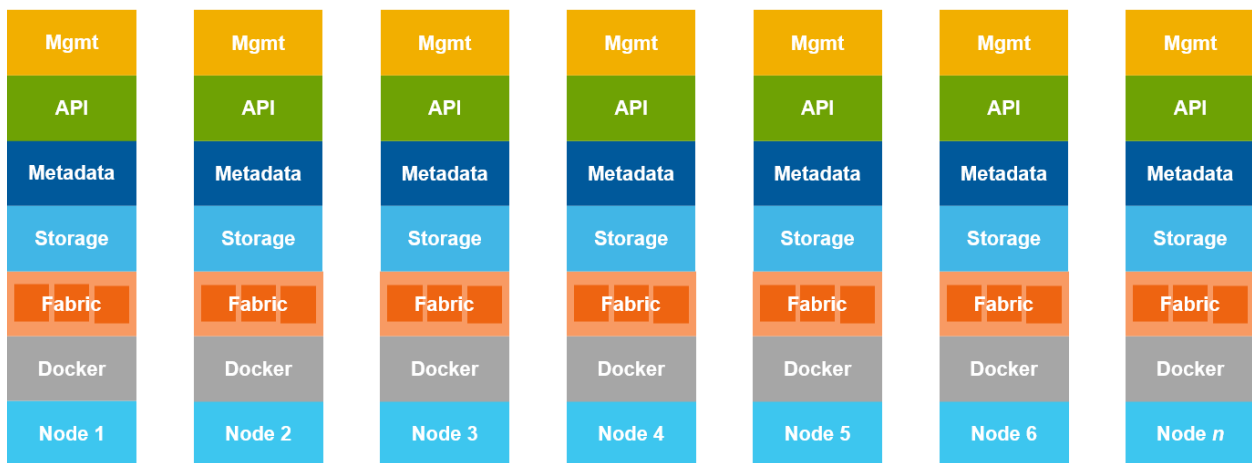
Figure 1: ECS Component layers (Ref.3)

- **Data Services:** The data services layer supports various object protocols; S3 (Simple Storage Service), Atmos, Swift and CAS (Content Addressable Storage) as well as file system protocols like HDFS and NFS (Network File System) for accessing the ECS object store. In general, ECS supports multiprotocol access, which means that data ingested via one protocol can be accessed via another. Data ingested using S3 can, for example, be modified using Swift, NFS, or HDFS.
- **Provisioning Portal / Mgmt:** The ECS portal is the centralized access point where you go to manage the ECS appliance. It provides options for both monitoring and operating (managing, licensing, and provisioning) ECS nodes.

The following detailed reporting options are available through the portal:

- Site, storage pool, node, and disc capacity consumption.
  - Various data storage parameters such as latency, throughput, transactions per second, and replication progress rate are all monitored.
  - System component data, such as node and disc recovery status and hardware and process health statistics for each node, which aids in the identification of performance and system bottlenecks.
- **Storage engine:** The storage engine layer provides an unstructured storage engine for data storage and retrieval, transaction management, and data protection and replication. Objects ingested via numerous object storage protocols, as well as the NFS and HDFS file protocols, are accessible through the storage engine.

- **Fabric:** Cluster health management, software management, configuration management, upgrade capabilities, and alerting are all provided by the fabric component layer. The fabric layer is in charge of keeping services up and running as well as managing resources like discs, containers, firewalls, and networks. It monitors and reacts to changes in the environment, such as failure detection, and sends out notifications about system health. Fabric Firewall Manager protects the 9069 and 9099 ports, which are public IP ports. Outside of the cluster, the port is unavailable.
- **Infrastructure:** ECS runs on top of SUSE Linux Enterprise Server 12 (SLES 12) operating system as a Java application; hence, Java Virtual Machine (JVM) is installed as well. Docker is installed on the infrastructure to deploy the encapsulated ECS layers. The containers are isolated only sharing the underlying operating system resources and hardware.
- **Hardware:** The Gen3 or EX-series appliance models of ECS are based on standard Dell servers and switches.



**Figure 2: ECS Appliance: All Services Run on Every Node (Ref.4)**

As shown in Figure 2, every single node runs every single service. This will simplify the deployment model in an ECS appliance. However, it is limited in terms of flexibility.



## 2.2.2 ECS Storage Structure

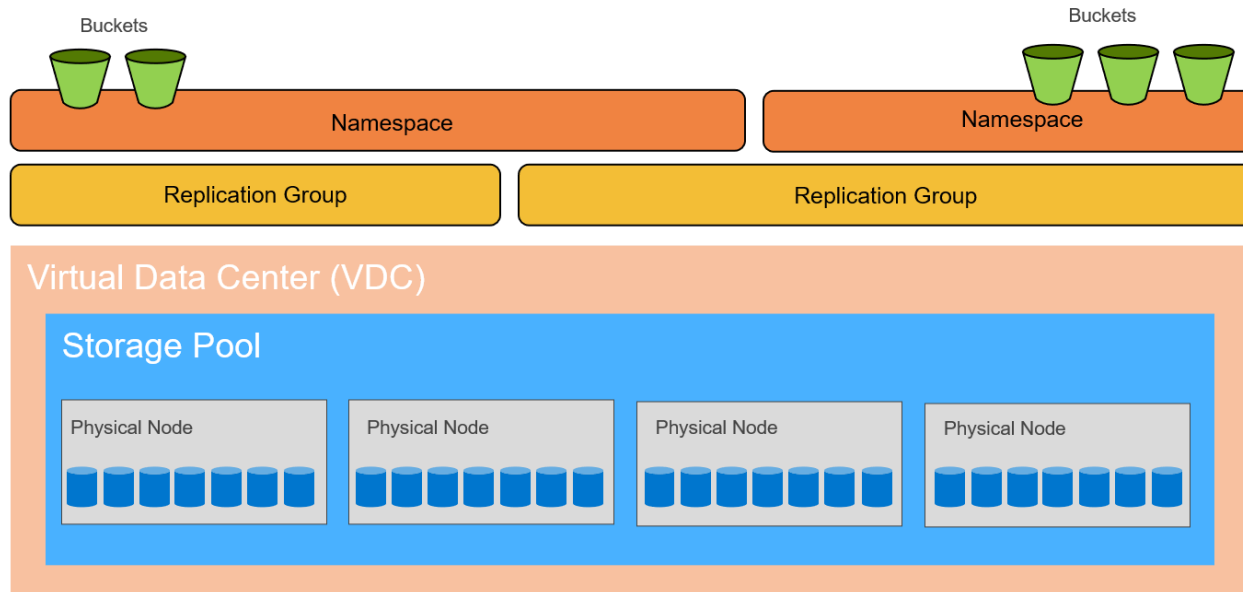


Figure 3 ECS Appliance: High-Level Design (Ref.4)

**Virtual Data Center** - A VDC is made up of a set of ECS infrastructure governed by a single fabric instance, generally referred to as a site or geographically distinct region.

**Storage Pool** - A storage pool (SP) is a subset of nodes and their associated storage that are part of a VDC. A storage pool can have any number of nodes; however, the minimum recommended is five. An ECS node can only belong to one storage pool. A storage pool is a tool that can be used to physically separate data from distinct applications. At the storage pool-level, erasure coding is done using either a 12+4 or a 10+2 scheme.

**Replication Group** - Replication groups (RG) specify where content from storage pools is safeguarded and where data can be read or written. Objects inside the same VDC are protected by local replication groups from disc, node, and rack failures. Global replication groups safeguard objects from disc, node, rack, and site failures across many VDCs. A VDC can be associated with numerous RGs.

**Namespace** - A namespace's most distinguishing feature is that users from one namespace are unable to access objects from another namespace. A namespace can represent a department or a group inside a department.

**Buckets** - Buckets are object data containers that are created in a namespace to allow apps to access data stored in ECS. These containers are referred to as "buckets" in S3, and ECS has embraced this term. In ECS, buckets are global resources. Each bucket has its own namespace, and each namespace has its own RG. A bucket is replicated across sites when the replication group spans multiple sites.

It is at the high-level design of the ECS appliance where we come in to dedicate a physical node and all its sub-resources. We allocate that into a storage pool. That storage pool is allocated to a VDC. Replication groups are sliced and behaviors for replication will be on top of the VDC and storage pools and so on. Lastly, we do our tenancy model across those buckets using Namespaces.

## 2.3 ObjectScale Architecture

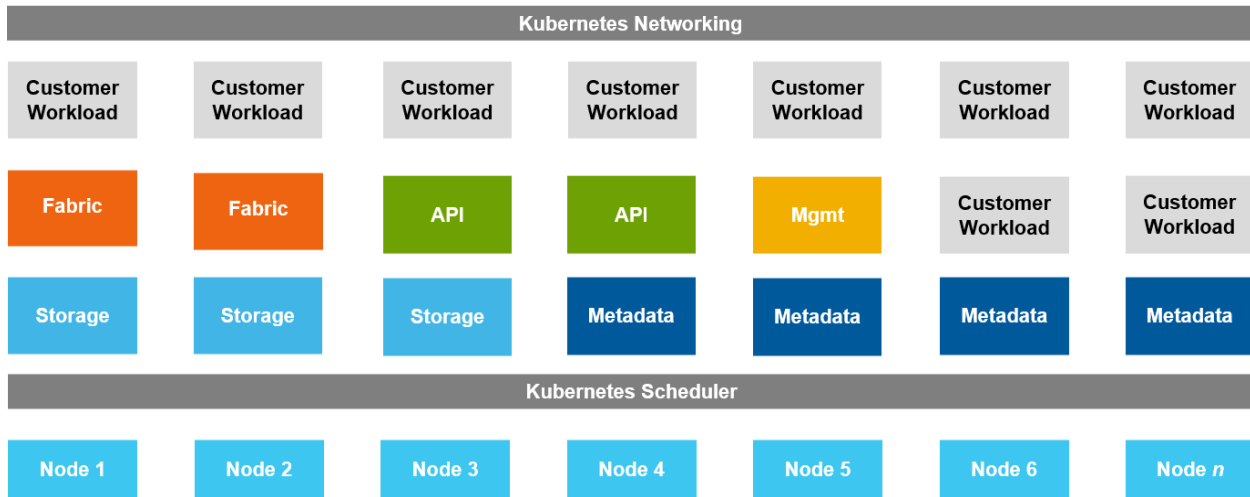


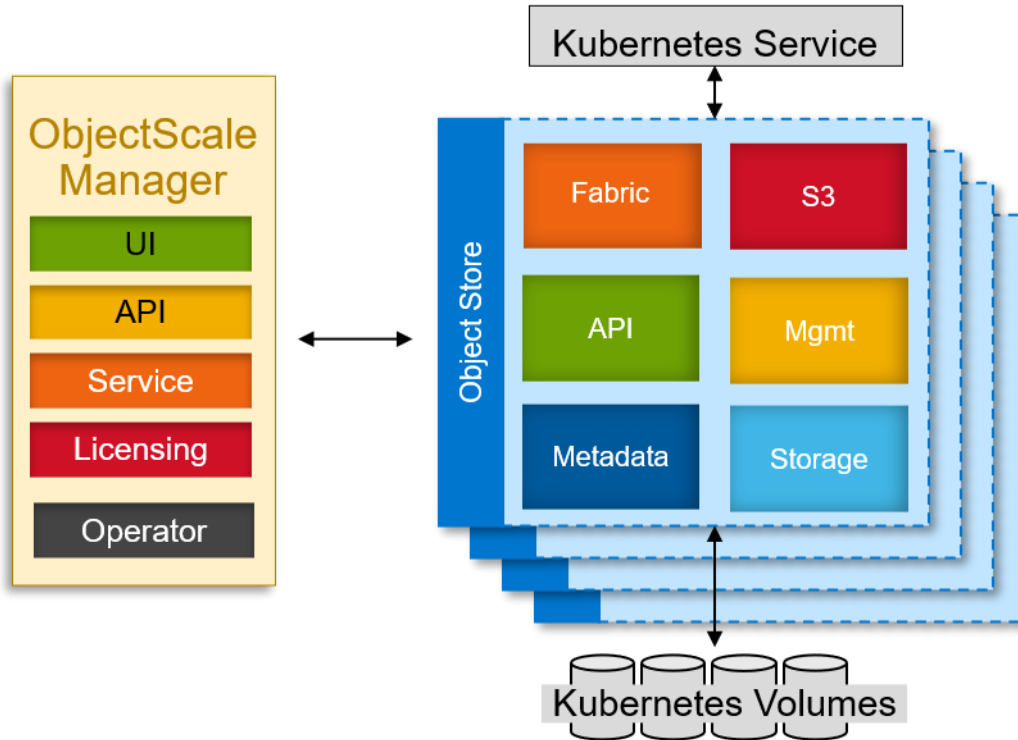
Figure 4: ObjectScale: Service Deployment (Ref.4)

ObjectScale takes advantage of Kubernetes' native orchestration features, such as scheduling, load balancing, self-healing, and resource optimization, to deliver enterprise-grade object storage in a simple, software-defined package.

ObjectScale uses the same Dell EMC ECS codebase. ObjectScale inherits the worldwide Global scalability, robust S3 compatibility, data protection, and security benefits that have helped make ECS the market leader in on-premises object storage. What makes it game-changing is that we do not have to map and use all of the underlying resources at any given time. ObjectScale enables distribution of the different storage services that interact with the physical disks, adding to the amount needed when it is needed. For example, choose the number of instances of the management services actually required. Another example would be having just two API gateways (Figure 4) versus having API gateways on every single node.

In short, ECS is basically monolithic containers, and with ObjectScale as architecture, ECS rips all that apart, micro-servicing it. So, we use what we need when we need it plus the added capabilities that come with using Kubernetes for orchestration such as load balancing, scheduling, self-healing, resource optimization, etc.

### 2.3.1 Object Store



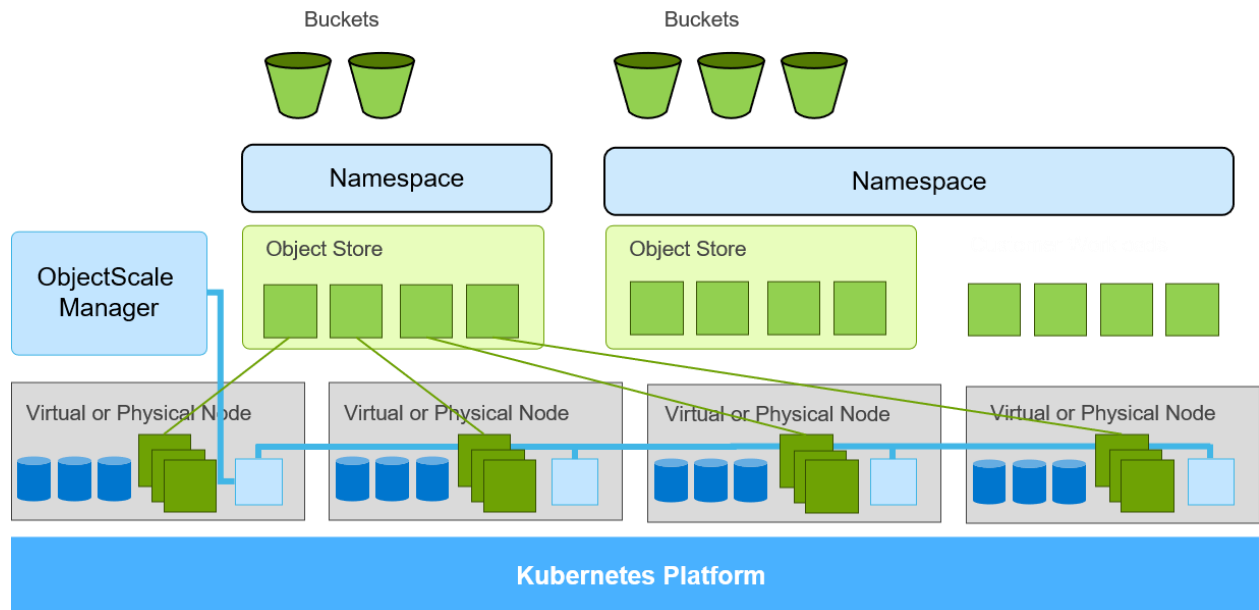
**Figure 5 ObjectScale: Object Store as a Kubernetes Application**

The next question is how does ObjectScale then differ in terms of architecture? Principles that drive ObjectScale from an infrastructure perspective is Kubernetes. With the ECS appliance, we use VDC, storage pools and such to drive object storage. With ObjectScale, we are moving into a construct called Object Store, a completely resource-isolated deployment of an object platform. ECS's processes are broken up into individually scaled Kubernetes "pods" to efficiently use only the resources needed. ObjectScale allows multiple object stores in one Kubernetes cluster. The object stores will interface with a global ObjectScale manager that orchestrates them: create, upgrade, delete etc. (lifecycle of multiple object stores).

A few services that are single per region are:

- **Identity and Access Management:** IAM is framework of policies and technologies that ensure the right users have the appropriate access to technology resources. ObjectScale supports a wide set of granular permissions, conditions, and policies. It enables bucket, object and user tagging and tags to be referenced in IAM policies.
- **Federation:** managing a geographically distributed environment as a single logical resource.
- **A web-based GUI** to manage, license and provision ObjectScale clusters.
- **GraphQL:** an open-source data query and manipulation language for APIs, and a runtime for fulfilling queries with existing data.

Each of the individual object stores will have different deployment profiles based on memory, CPU, and tasks. To optimally configure the storage environment for various workload demands, ObjectScale provides users an intelligent sizer tool. Simply provide ObjectScale with key parameters such as capacity and estimated data growth rate, and ObjectScale will provision the appropriate resources automatically.



**Figure 6 ObjectScale: High-Level Design (Ref.4)**

Disks will be provided into an object store and the object store will then be the instance of what we are doing with the deployment. Looking at a high-level design of objectscale, we have the Kubernetes platform, which has either virtual or physical nodes. We will also have our ObjectScale manager. Kubernetes acts as a connector between physical infrastructure, such as disc and network, and containerized application services. ObjectScale relies on Kubernetes' resource-management features to manage operating-system and hardware interactions.

The global infrastructure model for ObjectScale provides a regional construct where namespaces can be shown in one region or multiple regions. This enables replication of a bucket using Cross Region Replication (CRR) from the object store in one region to another bucket object store in another region (called as global replication) or even locally replicate the buckets. Replicating data across environments helps organizations protect workloads from outages and share data with teams across the globe, while native multi-tenancy provides resource isolation and secure global access through multi-site failure tolerance. ObjectScale Replication can be used to fuel everything from dev/test sandboxes to globally distributed data lakes.

### 3. Deployment Options with ObjectScale

ObjectScale runs on Kubernetes distributions including Red Hat OpenShift and VMware vSphere with Tanzu, allowing developers to focus on code while IT professionals manage ObjectScale Storage to policy. With deep VMware integration, ObjectScale can be deployed on Dell EMC VxRail, a hyper-converged infrastructure (HCI) system purpose-built for VMware environments. It's flexible IT at your fingertips. Let's discuss further.

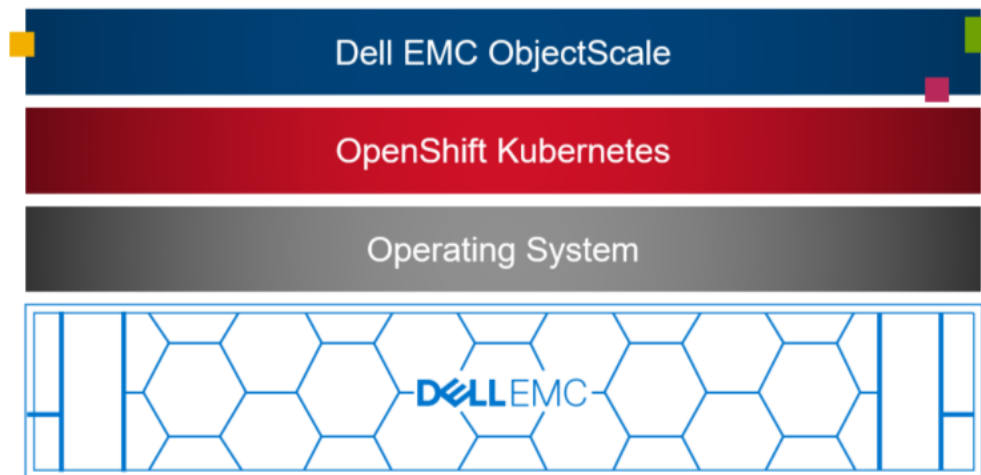
#### 3.1 ObjectScale on Red Hat OpenShift

The Red Hat OpenShift Container Platform is an on-premises private platform-as-a-service product. It helps enterprises establish a DevOps model for developing cloud-native apps by providing developers and operations teams with the tools and services they need to address today's needs while also planning for the future. It is built around application containers powered by Docker, with orchestration and management provided by Kubernetes, and runs on Red Hat Enterprise Linux and Container Linux.

Dell Technologies in association with Red Hat enables enterprises to deploy and run S3-compliant object storage on the OpenShift Container Platform to support cloud-native apps from an on-premises data center.

##### 3.1.2 Architecture

With ObjectScale's integration with Red Hat OpenShift, developers can use Kubernetes APIs to provide and scale high-performance object storage while providing IT with unrivalled simplicity, manageability, and reliability.



**Figure 7: Dell EMC ObjectScale on Red Hat OpenShift (Ref.6)**

Deploying ObjectScale with OpenShift alongside traditional applications in a full-stack infrastructure that comprises servers, primary storage, and supporting software is one approach to maximize business value.

Additional benefits of ObjectScale on OpenShift include:

- Designed for scalable data centers; clusters can scale up to petabytes in size and have 1000+ nodes.

- An ObjectScale bare-metal Container Storage Interface (CSI) driver and third-party container network interfaces (CNI) such as Calico and VMware NSX enhances storage efficiency, optimizes performance, and eases management.
- A certification tool exists to assist qualification of hardware for ObjectScale deployments.
- Long-term support for Red Hat OpenShift.

### 3.1.3 Dell Technologies Services for ObjectScale on Red Hat OpenShift

Customers can choose from a variety of options for ObjectScale support and deployment.

- **Dell EMC ProDeploy and ProDeploy Plus for Enterprise** offer remote and in-person deployment options along with site readiness review, implementation planning, etc.
- For qualifying 3rd party software, **Dell EMC ProSupport and ProSupport Plus for Enterprise** offers collaborative or single source support options.
- **Dell EMC Data Migration services** include a project manager and design specialists who use tools, techniques, and worldwide best practices to migrate data from Dell EMC and third-party storage to new infrastructure securely and efficiently.

### 3.2. Dell EMC ObjectScale on VMware vSphere with VMware Tanzu

VMware is dominating applications and infrastructure modernization. VMware created the ideal environment for running modern apps within the enterprise with the release of VMware Cloud Foundation with Tanzu in April of 2021. The latest developments in VMware vSphere and VMware vSAN, which foster collaboration between development and IT operations, are at the heart of VMware Cloud Foundation with Tanzu.

With Kubernetes integrated into vSphere, data center managers can manage both virtual machines and Kubernetes clusters from a single interface, on the same platform they already know and love. Developers may quickly deploy infrastructure services on demand thanks to an open ecosystem of Kubernetes-based APIs. Customers can now provision persistent cloud-scale object storage for modern stateful applications using VMware vSAN DPP (Data Persistence platform), which is integrated with VMware Cloud Foundation with Tanzu.

The vSAN DPP provides a platform for VMware partners who deliver modern stateful services to interact with the underlying virtual infrastructure, enabling execution of stateful services with high velocity scalability, simpler IT management, and lower TCO. VMware vSphere with Tanzu (VMware vSAN, VMware NSX, and VMware Tanzu) licenses are also required for the vSAN Data Persistence platform.

### 3.2.1 Architecture

Dell EMC ObjectScale's integration with VMware vSphere with Tanzu provides developers freedom to provision and scale high performance object storage via Kubernetes APIs while ensuring unmatched simplicity, manageability, and reliability for IT.

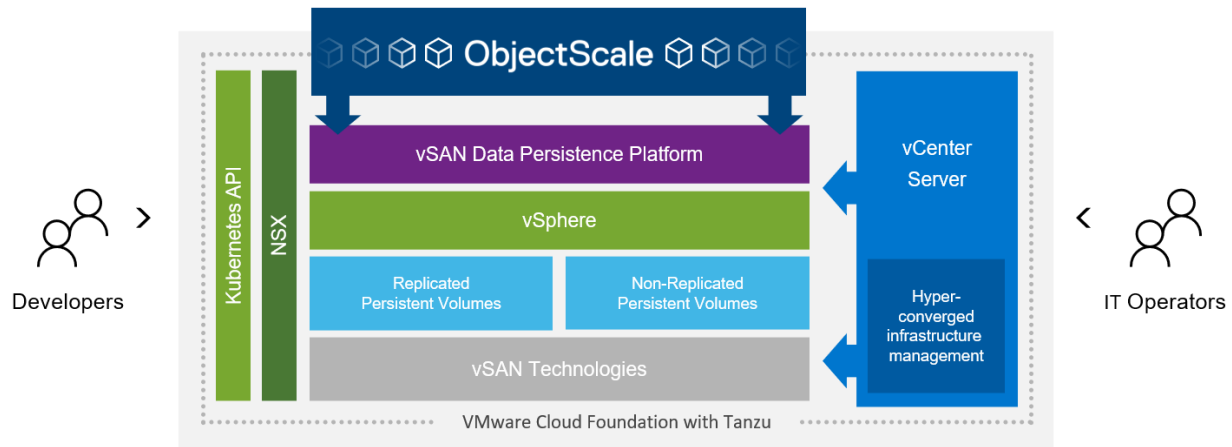


Figure 8: ObjectScale on VMware Cloud Foundation with Tanzu (Ref.8)

Organizations can use VMware vSAN Direct Configuration to deploy ObjectScale alongside traditional applications on a regular vSAN cluster using the vSAN-SNA (vSAN Support for Shared Nothing Architecture) policy or use VMware vSAN Direct Configuration (a technology enabling direct access to the underlying direct-attached hardware which can be optimized for the application needs) to deploy it on a dedicated vSAN cluster. Both choices benefit from excellent storage efficiency for stateful services by employing service-level replication and unified management of services in VMware vCenter.


### 3.3 Dell EMC ObjectScale on VMware vSphere with Tanzu on Dell EMC VxRail

Containers and Kubernetes are driving enterprise-wide modern application initiatives. While cloud-native applications promise to speed up development and innovation, many of these next-generation apps require modern S3-compatible storage to function properly. To fully embrace a genuine DevOps model, IT must also give developers with self-service infrastructure resources, such as storage, to enable smooth integration into CI/CD pipelines. To add to the complexity, IT departments are frequently required to support new apps and cloud-native efforts using current data center resources and skills.

#### 3.3.1 Advantages of ObjectScale on VxRail

Organizations will be able to support Kubernetes workloads and containers with ObjectScale software-defined object storage by running VMware vSphere with Tanzu and the vSAN Data Persistence platform on Dell EMC VxRail, the only fully integrated, pre-configured, and proven HCI system designed for VMware vSAN. It provides a turnkey solution for quicker, better, and simpler workload management in private and hybrid cloud settings.





**Figure 9: Dell EMC ObjectScale on Dell EMC VxRail (Ref.10)**

ObjectScale can be implemented on VxRail thanks to its integration with the vSAN Data Persistence technology, allowing organizations to take advantage of new modern application opportunities on familiar hyperconverged infrastructure.

**Powering new generation applications:** ObjectScale provides enterprise-grade storage on VxRail infrastructure, powering new modern apps that require S3 to run. S3 compatibility gives developers a familiar set of APIs to work with while modernizing existing workloads and designing the next generation of enterprise apps. With ObjectScale Lock (stores objects in a write-once-read-many model), IAM, Select (retrieves only the data needed for the application), Event Notifications (notifies users of bucket events), and ObjectScale Replication (bucket-to-bucket replication), ObjectScale delivers on the newest S3 APIs, allowing workloads developed for the cloud to run effortlessly on VxRail. Due to its performance, scalability, flexibility, and ease of use, ObjectScale is the ideal data store for new workloads such as big data analytics, IoT storage, media content distribution, artificial intelligence (AI), and machine learning (ML). ObjectScale can also be used as a low-cost secondary storage layer, freeing up expensive primary storage and enabling typical applications like backups, long-term storage and tape replacement.

**Consistent infrastructure and operations managed from VMware:** With ObjectScale's integration with VMware's vSAN Data Persistence platform, IT teams do not have to learn new skills or manage an external storage footprint. They can help developers with cloud-native storage using the VMware tools and capabilities they've developed in-house. They can now control shadow IT while freeing developers to use Kubernetes APIs to support CI/CD procedures and agile techniques in a self-service manner by running ObjectScale on VxRail.

Storage Policy-Based Management (SPBM) can be extended from VM volumes to container volumes, providing for a more uniform management strategy. Furthermore, VxRail Manager and vCenter are 100% natively integrated, and all VxRail management is done through the familiar vServer interface. Intelligent lifecycle management (LCM) automates non-disruptive updates, patches, node additions, and node retirements while keeping the VxRail infrastructure in a continuously validated state to guarantee workloads operate and clusters are optimized. ObjectScale and VxRail combine to provide persistent cloud-scale object storage for modern stateful applications on a uniform architecture with a consistent management experience, minimizing TCO.

**Storage infrastructure supported by a single vendor:** Storage supported by a single vendor - Dell Technologies: Companies can take advantage of the integrated storage stack offered by Dell with ObjectScale and VxRail, providing a consistent, unified experience for all storage support and maintenance requirements. Dell Technologies provides a comprehensive set of deployment, support, and advisory services to help you get the most out of your technology. Dell EMC ProDeploy, ProSupport, and Data Migration are some of the services available. To comply with the various financial strategies that businesses follow, flexible consumption models and payment solutions are also available.

## 4. ObjectScale Use Cases

**Cloud-native web and mobile apps:** ObjectScale is designed specifically for cloud-native apps. Developers may incorporate object stores into their CI/CD (Continuous Integration / Continuous Delivery) pipelines in a familiar, self-service manner, expediting application development, thanks to strong S3 compatibility for data access and Kubernetes APIs for container management. You can simply keep up with data expansion as applications grow in scope since you can scale without boundaries. Data can be dispersed over any number of sites with ObjectScale Replication to ensure fault tolerance and low latency access.

**Dev / Test:** ObjectScale is the appropriate sandbox for S3-enabled modern application development. Object stores and buckets are simple to install thanks to automatic provisioning and an integrated workload sizer. Identity and access management (IAM) controls offer safe data access, while multi-tenancy enables resource separation. You can set up scratch object stores in minutes, decommission them when they're no longer needed, and then push them to production with a few simple clicks using ObjectScale. Data can be easily shared between development teams, allowing for more collaboration and productivity.

**Consolidated data lake:** Sensor telemetry, machine-generated logs, and application data may all be stored in an indefinitely scalable, centralized data lake using ObjectScale. You can choose the performance profile that best fits your workloads by deploying it on all-flash or capacity-optimized media. Connect several sites through federation to avoid data silos and offer data access from the edge to the core. Tag objects to improve analysis and make data more discoverable. Cloudera, Dremio, and Vertica provide powerful data lake engines and analytics systems.

**AI & ML:** ObjectScale supports next-generation AI and ML workloads, the next frontier for object storage, by delivering performance at scale. ObjectScale offers large datasets at high transfer rates to the most demanding CPU and GPU servers, allowing AI training algorithms to access more data without the need for HPC storage. Clusters may be easily scaled out to increase performance and capacity in a linear fashion. Object tagging also gives inference models more data to work with, allowing them to generate better predictions.

**Analytics:** Generate operational insights at the pace your organization requires by running lightning-fast queries on a performant ObjectScale data lake. Storage performance is no longer a limitation thanks to the option to deploy on NVMe-based all-flash devices. Custom tags give data more context, allowing for easier discovery and faster results. S3a allows Hadoop applications to read and write data to ObjectScale without having to manage a sophisticated HDFS cluster.

**Backup and archive:** ObjectScale is an S3 backup target and long-term archive with low TCO. ObjectScale protects data from ransomware attacks and assures robustness from node or disc failures with ObjectScale Lock for data immutability, data-at-rest-encryption (D@RE), global replication, and erasure coding protection techniques. Everything from virtual machines to Microsoft Office 365 backups is safe and accessible with ObjectScale. It's scaled-up business continuity.

**Enhancing Dell EMC VxRail:** ObjectScale adds S3-compatible, cloud-native object storage to Dell EMC VxRail, allowing it to serve workloads ranging from high-performance analytics to long-term archiving. ObjectScale can handle stateful sets thanks to the VMware vSAN Data Persistence platform, which can be configured directly through the vCenter interface. You can now support modern application development projects without having to acquire new infrastructure because VxRail is integrated, optimized, tested, and verified to run VMware's stack.

## 5. ObjectScale and ECS Positioning

To position ObjectScale and ECS, it is important that we understand their differences.

Parameters	ObjectScale	ECS
Appliance	Future.	Yes. HDD and NVMe SSD.
Software-defined	Yes. Containerized / microservices architecture orchestrated via K8s.	Yes. Runs on certified hardware.
Architecture	True microservices. Each service is deployed as a microservice.	Monolithic. Although, containerized and deployable via Docker.
Data Services	HTTP (S3)	HTTP (S3), NFS, SWIFT, ATMOS, CAS (others through gateways: CIFS, SMB, etc.)
Consistency model	Local buckets that are globally replicated with asynchronous replication. ObjectScale is strongly read-after-write within an Object Store.	Strong consistency globally and locally.
Global Redundancy	Globally redundant through ObjectScale Replication (OSR) with no limits on the number of sites.	Replication groups (RG) capable of spanning up to 8 virtual data centers (VDC).
Multi-tenancy	Isolate resources, shared and dedicated tenancy.	Storage pool would become dedicated resources.
Serviceability	Customer-deployable, customer-upgrade, self-service, professional services.	Professional Services for upgrades, migrations.
Licensing	Subscription, perpetual.	Perpetual.

**Table 1: Differences between ObjectScale and ECS (Ref.11)**

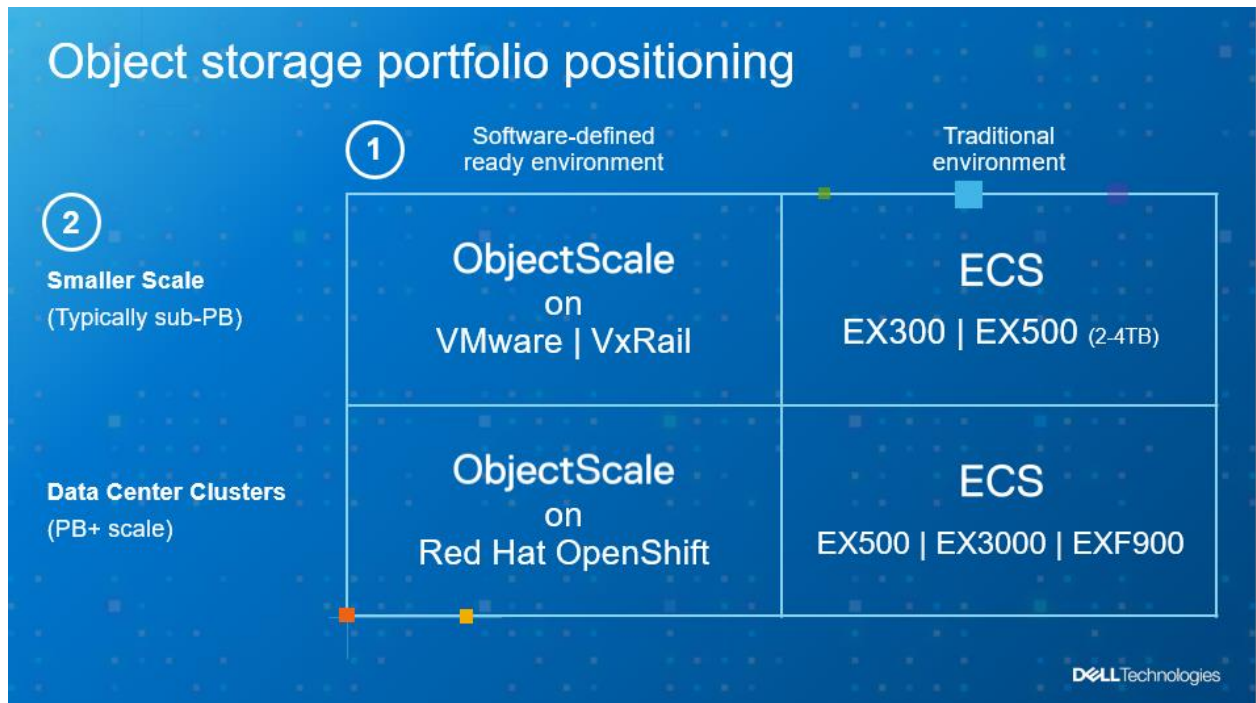


Figure 10 Dell Technologies' Object storage portfolio positioning (Ref.11)

ECS is Dell EMC's current enterprise-grade object storage system that features a microservices framework. ECS is designed for organizations that require a fully integrated turnkey appliance for diverse object storage requirements, though it can also be deployed as a software-defined solution that sits on top of certified hardware. ECS has a long legacy of enterprise-proven functionality and performance.

ObjectScale uses functionality native to the Kubernetes platform, with architecture that can be easily deployed and scaled. ObjectScale builds upon Dell's success in the object storage market by implementing a software-defined stack that enables object storage capabilities for VMware and Red Hat today, and additional use cases and platforms in the future.

## 6. Conclusion

ObjectScale clusters, which are built on a scale-out architecture, may extend from a few terabytes to petabytes and beyond, with no restriction on the amount of object stores, buckets, or objects stored. Clusters can begin with just three nodes and expand from there. ObjectScale Replication allows you to replicate object data from the edge to the core data center, from any location where an ObjectScale footprint exists. By replicating data throughout the environment, users may protect workloads against failures and share data with colleagues worldwide. D@RE, erasure coding, versioning, resource isolation, IAM, WORM (Write Once Read Many) features, and more are included in ObjectScale's enterprise security and data protection capabilities. It's storage technology you can rely on from the industry leader.

Having two product offerings in ECS and ObjectScale enable Dell Technologies to better meet the increased demand for object storage while serving two unique use cases and deployment models. Both products are an important part of the Dell Technologies product portfolio and both are designed to provide enterprise-grade object storage. However, each is optimized for different usage and deployment scenarios.

ECS will continue to serve organizations that require large-scale industrial archive and enterprise solutions using some of the densest appliance form factors in the market.

ObjectScale is ideal for edge use cases or when organizations want to deploy flexible, lightweight object storage close to the applications they support. Due to its software-defined nature and S3 compatibility, ObjectScale can also be deployed on existing virtualized infrastructure, making it an excellent platform for cloud-native applications and dev/test scenarios.

## 7. References

1. <https://www.gartner.com/doc/reprints?id=1-27L5HBBA&ct=211004&st=sb>
2. <https://www.redhat.com/en/topics/data-storage/software-defined-storage>
3. <https://www.delltechnologies.com/asset/en-us/products/storage/industry-market/h14071-ecs-architectural-guide-wp.pdf>
4. <https://inside.dell.com/videos/84642>
5. <https://www.delltechnologies.com/asset/en-in/products/storage/briefs-summaries/h18906-dell-emc-solution-overview-objectscale.pdf>
6. <https://www.delltechnologies.com/asset/en-in/products/storage/briefs-summaries/h18908-dell-emc-solution-overview-objectscale-on-red-hat-openshift.pdf>
7. <https://en.wikipedia.org/wiki/OpenShift>
8. <https://www.delltechnologies.com/asset/en-in/products/storage/briefs-summaries/h18907-dell-emc-solution-overview-objectscale-on-vcf-with-tanzu.pdf>
9. <https://www.dell.com/en-us/blog/dell-emc-objectscale-early-access-on-vmware-cloud-foundation/>
10. <https://www.delltechnologies.com/asset/en-in/products/storage/briefs-summaries/h18905-dell-emc-solution-overview-objectscale-on-vcf-with-tanzu-on-vxrail.pdf>
11. <https://inside.dell.com/docs/DOC-515480>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.