



Hewlett Packard
Enterprise

DriveScale-HPE Reference Architecture

Table of Contents

Executive Summary	3
Audience and Scope	3
DriveScale Advantage	3
HPE Advantage	4
HPE or Industry standard JBOD	4
DriveScale HPE Solution Overview	5
Entry or POC Configuration	5
Full Rack Configuration	6
Storage Rack Configuration	7
Rack Scalability	7
Summary	8
Conclusion	8

DRIVESCALE-HPE

Executive Summary

DriveScale has engineered the next-generation Software Composable Infrastructure (SCI) that is designed to fundamentally change the way data center architects design, deploy, manage and consume the hardware and software resources. DriveScale provides capabilities to IT operators to connect disaggregated resources in an intelligent manner, the ability to manage, modify and scale these components over time. SCI results in a simpler deployment model with a fluid pool of resources that can be used for modern Big Data workloads with significantly improved agility.

This document is a high-level design reference architecture guide for implementing a DriveScale solution with HPE ProLiant or Apollo rack-optimized or Blade servers and HPE or industry standard JBODs for modern workloads such as Big Data.

The reference architecture introduces all the high-level hardware and software component that are included in the stack. Each high-level component is then described individually.

Audience and Scope

This reference architecture guide is for modern workloads such as Big Data workloads and IT architects who are responsible for the design and deployment of modern workloads on premises, as well as for Big Data administrators and architects and data center architects/engineers who collaborate with specialists in that space.

DriveScale Advantage

Activate your HPE Rack optimized datacenter with DriveScale for modern workloads

Modern workloads such as Big Data workloads have become integral part of the traditional data centers. Typical Big Data deployments in any enterprise start with small and isolated deployments. They are designed to scale by adding more hardware resources to the same cluster. This flexibility is built into the Big Data software.

However, typical Big Data deployments have multiple limitations:

- Administrators can't respond quickly to changing application stacks and data velocity.
- Deployments are over-provisioned with resources and under-utilized in order to provide service level guarantees.
- Multiple silos of hardware are created for each application workload.

Once installed, the DriveScale Management System automatically discovers the installed components and their network connectivity and inventories them. It also discovers network bandwidth domains to ensure that drives in the software composed servers are within the same bandwidth domain for optimal performance. It then uses this inventory to compose logical servers in software using either Restful API that can be interfaced to customer's management system, or use DriveScale GUI.

DriveScale software composed servers deliver nearly the same performance as custom configured hardware servers in all Big Data micro-benchmarks, and match the performance in real Big Data application loads.

DriveScale Management System enables a customer to define alarms on many hardware metrics to enables alerts to be generated when these alarms are triggered. The DMS enables a customer to diagnose the system at logical server level or physical component level.

DriveScale’s rack scale architecture provides the solutions to all the modern Big Data workload deployment limitations. **Now administrators can flexibly deploy and manage independent pools of resources at a lower cost without making changes to the application stack.**

HPE Advantage

HPE offers the world’s most secure industry standard servers. HPE’s approach to excellence in the innovation and quality that is instilled across the product lifecycle, from the customers first approach to design, to the supplier selection, quality and management, to the world-class manufacturing and rigorous product testing, to HPE’s global support services and network of channel partners. The HPE servers aspires to make IT simpler.

- Makes your organization agile
- Protect your digital assets and data with unique security
- Control your financial destiny in IT

HPE or Industry standard JBOD

With DriveScale solution, we recommend customers should use high capacity JBODs with dual hot-pluggable IO controllers (Expanders) and enough upstream bandwidth. The JBODs should also have dual hot-pluggable redundant power supplies. DriveScale has evaluated and tested a few of the vendor offerings for redundancy, management functionality and performance. The table listed below has the HPE supported JBOD’s along with other supported industry vendors name and the model numbers which are certified by DriveScale.

Table 2: DriveScale certified JBOD

JBOD Vendor	Model Number
Hewlett Packard Enterprise	D6020 - 3.5", 70 bays, 5U, quad expanders, 4 x 2 x mini-SAS 12G
	D6000 - 3.5", 70 bays, 5U, quad expanders, 4 x 2 x mini-SAS 6G
RAID Inc./Newisys	NDS-4600/4603 - 3.5", 60 bays, 4U, redundant expanders, 2 x 4 x mini-SAS 6G
	NDS-2241 - 2.5", 24 bays, 2U, redundant expanders, 2 x 3 x mini-SAS 6G
	NDS-4900 - 3.5", 90/96 bays, 4U, redundant expanders, 2 x 6 x mini-SAS-HD 12G
	NDS-4900 - 3.5", 84 bays, 4U, redundant expanders, 2 x 5 x mini-SAS-HD 12G
Qunta (QCT)	M6400H - 3.5", 60 bays, 4U, redundant expanders, 2 x 4 x mini-SAS 6G
	JB4602 - 3.5", 60 bays, 4U, redundant expanders, 2 x 4 x mini-SAS 12G

JBOD Vendor	Model Number
Promise Inc.	J5300s - 3.5", 12 bays, 2U, redundant expanders, 2 x 2 x mini-SAS-HD 12G
	J5320s - 2.5", 24 bays, 2U, redundant expanders, 2 x 2 x mini-SAS-HD 12G
	J5600 - 3.5", 16 bays, 3U, redundant expanders, 2 x 2 x mini-SAS-HD 12G
	J5800 - 3.5", 24 bays, 4U, redundant expanders, 2 x 2 x mini-SAS-HD 12G
Dell	PowerVault MD3060e - 3.5" and 2.5", 60 bays, 4U, redundant expanders, 2 x 3 x mini-SAS 6G
Lenovo	D3284 - 3.5", 84 bays, 5U, redundant expanders, 3 x 4 x mini-SAS-DS 12G

DriveScale | HPE Solution Overview

Entry or POC Configuration

The reference architectures shown below describes a typical single-rack configuration for deploying a Big Data infrastructure utilizing HPE ProLiant servers, HPE networking ethernet switches and HPE standard direct-attached storage in the form of JBOD's.

The first design shown is a suggested entry configuration that might be deployed for a small workload or for proof-of-concept testing. This design highlights the value proposition of the DriveScale/HPE solution in that it demonstrates how customers can achieve significantly higher levels of **flexibility** in deployment, **reduction in cost** by starting small and **seamless integration** with Big Data workloads such as Hadoop with a highly software composable solution.

The entry or POC configuration is based 4 HPE DL 360 disk lite servers serving as the data nodes and 1 HPE DL 360 disk lite server as the master node. The storage for the Big Data workloads are provided by the 70 drives hosted in the HPE JBOD. Using the DriveScale Adapter and DMS, customers can mount drives to each of the nodes. They can also expand the capacity of the cluster by adding more nodes if they application requires more compute or add more drives to the existing nodes from the JBOD. Customers have the flexibility to spread the 70 drives across the 5 nodes or add more nodes to be used with the existing 70 drives in the JBOD.

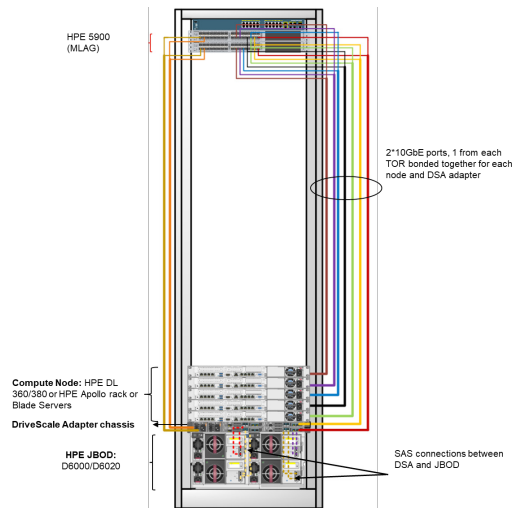


Figure 1: POC/Start up config

Full Rack Configuration

The design shown in this section builds on the first reference design shown above, to show what a full rack might look like. In this architecture, the aim is to achieve the highest density for servers by using the HPE ProLiant or Apollo rack or blade servers, HPE JBOD with high bandwidth HPE networking and storage. The design also takes into consideration the optimal ratio of drives to servers and the bandwidth to each drive, to build out a rack-level configuration that is optimized for the typical big data workloads. Scaling is achieved by simply replicating this rack-level configuration. Customers have the flexibility to spread the 280 drives across the 15 nodes or add more nodes to be used with the existing 280 drives or have less than 15 nodes with 280 drives in the JBOD. Customers can work with DriveScale Solution Architects to optimally grow the rack and cluster.



Figure 2: Full rack config

Storage Rack Configuration

The design shown in this section builds on the first reference design shown above, to show what a full rack might look like. In this architecture, the aim is to achieve the highest density for servers by using the HPE ProLiant or Apollo rack or blade servers, HPE JBOD with high bandwidth HPE networking and storage. The design also takes into consideration the optimal ratio of drives to servers and the bandwidth to each drive, to build out a rack-level configuration that is optimized for the typical big data workloads. Scaling is achieved by simply replicating this rack-level configuration. Customers have the flexibility to spread the 350 drives across the 9 nodes. Customers can work with DriveScale Solution Architects to optimally grow the rack and cluster.



Figure 3: Storage rack config

Rack Scalability

This design also demonstrates that one can start small and add components as needed, to scale the system. Customers can grow their compute and storage depending on the application needs. If the application demands more compute, then they can add HPE ProLiant or Apollo rack or Blade servers and if the application demands more storage then they add HPE or any JBOD's to the cluster. They can scale from few nodes and JBOD's to thousands of nodes and JBOD's.




Figure 3: Rack scalability

Summary

In all the above reference designs, HPE servers, DriveScale Adapters and HPE JBOD's are co-located in the same rack. This ensures that drives are always only one Ethernet switch hop away from the servers that they are attached to, ensuring optimal performance. DriveScale's solution is designed to maintain the data locality that Big Data applications such as Hadoop requires and provide performance to disk that is equivalent to direct-attached drives.

Therefore, so long as we design the systems with no more latency in the network than that produced from a single Ethernet switch, the systems will perform optimally.

Given the composable nature of the DriveScale/HPE solution, one can scale as needed on either the compute elements or the storage elements very easily, thereby adding capacity as required. This gives the solution a very flexible architecture, not just in terms of composability but also in terms of capacity scaling. Customers can consult with the DriveScale Solution Architects to optimally design and deploy the Big Data workload hardware infrastructures.

Conclusion

The DriveScale/HPE reference architecture described in this document will run any standard Big Data workload such as Hadoop deployment including Cloudera, Hortonworks MAPR and Apache Hadoop. DriveScale solution architects are available to help customers design their infrastructure for optimal cost and performance.

With DriveScale, customers can now operate their on-prem datacenter with the flexibility of a cloud environment while spending less on infrastructure and operations. The architecture also integrates quickly and easily into the existing environment with no additional changes to the application stack.



DriveScale, Inc
1230 Midas Way, Suite 210
Sunnyvale, CA 94085

Main: +1 (408) 849-4651
www.drivescale.com