RED HAT®
HYPERCONVERGED
INFRASTRUCTURE

RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION

Tested and validated server configurations

DATASHEET

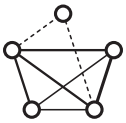

Easily virtualize your business applications, maximizing resource utilization.

Manage integrated compute-plus-storage resources with a simplified single management interface.

Increase productivity by focusing on applications instead of infrastructure.

Reduce risk and achieve predictable performance with tested and validated clusters of industry-standard servers.

Hewlett Packard Enterprise
Server Tested

facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

redhat.com

INTRODUCTION

Validated server configurations for Red Hat® Hyperconverged Infrastructure simplify and reduce the risk of designing and deploying hyperconverged computing infrastructure (HCI). Tested and optimized, these validated server configurations combine with Red Hat Hyperconverged Infrastructure for Virtualization to yield platforms that are durable and highly available. The validated server configurations help predictably consolidate infrastructure, producing operational efficiencies through optimized management workflows. In addition, unified storage and compute resources provide simplified, low-touch operations.

Red Hat Hyperconverged Infrastructure for Virtualization offers an open, simple, and optimized platform for your application workloads in a small footprint. Integrating Red Hat Virtualization and Red Hat Gluster® Storage, the solution also incorporates Red Hat Ansible® Automation, offering unified administration from a single platform. By eliminating the need for a discrete storage tier, the solution removes many of the traditional burdens associated with acquisition, set-up, and day-to-day operations—letting you focus on more business-critical tasks.

VALIDATED INDUSTRY-STANDARD CONFIGURATIONS

To evaluate performance and validate configurations, Red Hat conducted performance testing on HPE DL360 Gen10 and DL380 Gen10 servers, configured as shown in Table 1. Both capacity-optimized and throughput-optimized servers were evaluated, with the anticipation that similarly configured industry-standard servers from other vendors would yield similar results. Three-node Red Hat Hyperconverged Infrastructure for Virtualization clusters were then assembled, tested, and validated with representative workloads.[1]

TABLE 1. VALIDATED SERVERS FOR RED HAT HYPERCONVERGED INFRASTRUCTURE FOR VIRTUALIZATION

	CAPACITY-OPTIMIZED	THROUGHPUT-OPTIMIZED
Server platform	HPE DL360 Gen10	HPE DL380 Gen10
Minimum/maximum cluster size	3 servers / 12 servers	
Dual-socket processor (tested)	Intel Xeon Silver 4116 (12 cores per socket)	Intel Xeon Gold 6130 (16 cores per socket)
Memory/RAM (tested/max)	128GB/256GB	512GB/512GB
Network adaptor	HPE Ethernet 10Gb 2-port 562FLR-SFP+	HPE Ethernet 10/25Gb 2-port 640FLR-SFP28

1 *Independent testing, Red Hat Bangalore laboratory, September to November 2018.*

	CAPACITY-OPTIMIZED	THROUGHPUT-OPTIMIZED
Input/output (I/O) controller	HPE Smart Array P816i-a SR Gen10 (16 internal lanes, 4GB cache/ SmartCache) 12G SAS Modular Controller	
Data drives	HPE 2.4TB SAS 12G Enterprise 10K SFF (2.5in) SC 3-year warranty, 512e digitally signed firmware hard disk drive	HPE 6TB SAS 12G Midline 7.2K LFF (3.5in) SC 1-year warranty, 512e digitally signed firmware HDD
Data drive quantity	8	12
Operating system (OS) drives	HPE 960GB SATA 6G Read Intensive SFF (2.5in) SC 3-year warranty digitally signed firmware solid-state drive (SSD)	
OS drive quantity and data protection	2 (RAID 1)	2 (RAID 1)
Virtual machine (VM) density (3-node cluster example*)	42	90
Maximum ingest throughput (3-node cluster example**)	320MiB/s	705MiB/s

* *VM density calculated assuming average of four virtual central processing units (vCPUs) and 16GiB RAM per VM, with 300% CPU oversubscription and 150% memory oversubscription. This example also assumes that all VMs are high availability. In other words, all VMs can continue to run following the loss of one node in a three-node cluster.*

** *Ingest throughput was measured at 90/10 write/read I/O ratio via the fio open source benchmark utility.*

WORKLOAD CATEGORIES

Red Hat tested and validated server configurations for Red Hat Hyperconverged Infrastructure for Virtualization for the following workload categories:

- **General server consolidation**. Organizations increasingly want to consolidate general servers and proprietary storage appliances—particularly at edge locations—into a single hyperconverged server cluster. This approach eliminates special-purpose storage appliances and bare-metal servers, while providing flexible, highly available virtual machines (VMs) with corresponding data protection. This workload category does not typically exhibit particularly demanding input/output (I/O) characteristics, instead requiring only basic I/O performance with adequate storage capacity.
- **Log ingest**. This workload category ingests a steady stream of data into hyperconverged clusters at edge locations. Operators deploying this workload category typically seek to capture data from remote sensors and data acquisition equipment at the rate of data generation, while protecting against data loss through highly available VMs and protected storage. This workload category typically has high I/O throughput needs.

SERVER TESTING

Red Hat used the fio[2] and the DVD Store 3[3] tools to test the capabilities of the three-node hyperconverged clusters. Tests measured both I/O throughput and latency at increasing workload scales. All tests were performed from VM clients running within the Red Hat Hyperconverged Infrastructure for Virtualization cluster, transacting I/O with virtual disks backed by Quick Emulator (QEMU) Copy on Write (qcow2) images stored within the hyperconverged Gluster trusted storage pool. Dual 10GbE network connections ran between cluster nodes, one for front-end VM access and one for back-end VM management, migration, and storage traffic.

Throughput tests were performed with the fio tool as 30-minute timed tests using 4GiB file sizes, 4MiB block sizes, a 90%/10% write/read mix, and the direct=1 setting. Each Red Hat Hyperconverged Infrastructure for Virtualization host ran one load client VM and each client executed one fio job. Total write throughput was calculated as the sum of the three concurrent client jobs. Read throughput results were discarded. Following a warm-up test run, final results were then calculated from an average of three test runs.

CONCLUSION

Many organizations appreciate the benefits of hyperconverged infrastructure but require easier ways to purchase and deploy systems. Validated server configurations for Red Hat Hyperconverged Infrastructure for Virtualization help reduce the risk of deploying HCI. With tested and validated configurations for capacity-optimized and throughput-optimized workloads, organizations can consolidate infrastructure, easily virtualize applications, and increase productivity for IT staff.

2 freshmeat.sourceforge.net/projects/fio
3 github.com/dvdstore/ds3

ABOUT RED HAT

Red Hat is the world's leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.

facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

NORTH AMERICA
1 888 REDHAT1

EUROPE, MIDDLE EAST, AND AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com

redhat.com
#F14934_1118