

PERFORMANCE-INTENSIVE WORKLOADS WITH RED HAT STORAGE AND SAMSUNG NVMe SSDs

A solution for use cases previously out of reach of traditional OpenStack deployments

SOLUTION BRIEF

INTRODUCTION

Red Hat Storage and Samsung have partnered to create a solution that addresses several use cases that were previously out of the reach of traditional OpenStack® deployments, including:

- Fast pool of storage for private or public cloud service providers.
- Analytics workloads.
- Multiple distributed MySQL/MariaDB databases established as Database-as-a-Service (DBaaS).
- Network digital video recorder (DVR) quick retrieval of long-tail content.
- Telco edge network services.
- Financial service workloads.
- Test or development environments that must be staged or torn down quickly.

THE PROBLEM

In the past, it has been difficult to optimize the performance of OpenStack clouds in certain use cases. Red Hat® Ceph Storage has long been the de facto standard for creating OpenStack cloud solutions across block and object storage as a capacity tier based on traditional hard disk drives (HDDs). Now a performance tier using Red Hat Ceph Storage clusters and nonvolatile memory express (NVMe) solid state drives (SSDs) can be deployed in OpenStack environments. This performance tier is optimized to support the bandwidth, latency, and input/output per second (IOPS) requirements of high-performance workloads and such use cases as distributed MySQL databases, telco network DVR (NDVR) long-tail content retrieval, and financial services.

The Samsung NVMe Reference System is engineered to provide a well-balanced storage server node that includes matching CPUs, networking, storage, and peripheral component interconnect express (PCIe) connectivity to both deploy large amounts of NVMe SSDs and maximize Ceph's performance. The Ceph reference architecture can deliver an industry-leading 690K IOPS to IO-intensive workloads and 28.5Gb/s network throughput on a three-node cluster.

As a result, the reference architecture is an optimized pool of high-speed storage designed for OpenStack deployments, virtual infrastructures, financial service providers, and private and public clouds. In addition, the reference architecture can increase storage efficiency in test or development environments that must be deployed and dismantled quickly.

RED HAT AND SAMSUNG APPROACH

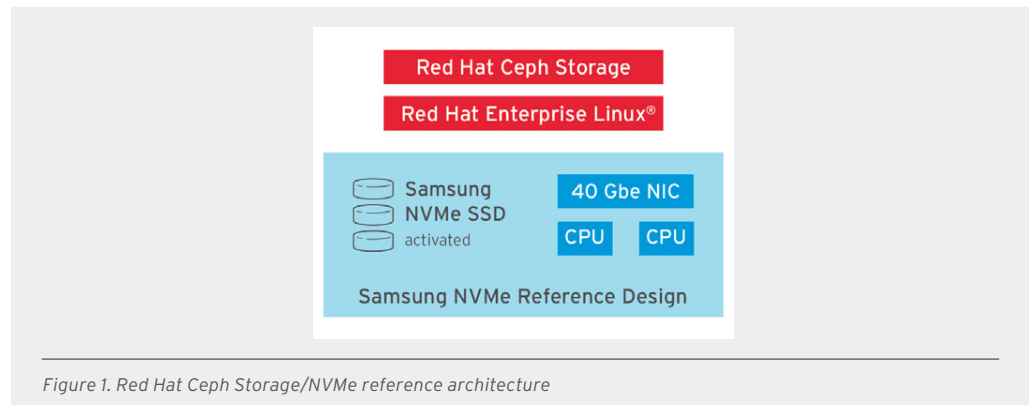
CEPH

Ceph is a key, established, open source software technology for scale-out, capacity-based storage under OpenStack. Ceph provides block-level, object, and file-based storage access to clusters based on industry-standard servers. Now Ceph supports a performance-optimized storage cluster using high-performance Samsung NVMe SSDs deployed in a Samsung NVMe reference design.



facebook.com/redhatinc
@redhatnews

linkedin.com/company/red-hat



RED HAT CEPH STORAGE

Red Hat Ceph Storage is a massively scalable, open source, software-defined storage system that supports unified storage for cloud environments. With object and block storage in one platform, Red Hat Ceph Storage efficiently and automatically manages the petabytes of data needed to run businesses facing massive data growth.

Red Hat uses the global open source community, including its own engineering resources, to develop new features and then locks down changes for predictable and stable releases. Red Hat Ceph Storage helps businesses automatically and cost-effectively manage their storage requirements, enabling those enterprises to focus on their own business needs and not their underlying IT infrastructures. As a software-defined storage platform, Red Hat Ceph Storage scales across physical, virtual, and cloud resources, giving organizations the ability to add capacity as needed without sacrificing performance or forcing vendor lock-in.

SAMSUNG NVMe SSDs

Samsung Enterprise SSDs are increasingly being used as data storage media in computing, communication, and multimedia devices. Most SSDs use NAND flash memory, which can retain data without an external power supply.

SSDs offer superior reliability compared to traditional HDDs. Advances in semiconductor flash memory technologies have allowed the development of SSDs that have much larger capacities than HDDs and can be used as direct replacements. SSDs also prove to be highly cost-effective due to their much lower power consumption and maintenance costs.

As the world leader in semiconductor memory technology, Samsung revolutionized the storage industry by shifting the planar NAND to a vertical structure. Samsung V-NAND technology features a unique design that stacks 48 layers on top of one another instead of trying to decrease the cells' pitch size. Samsung offers a comprehensive range of SSDs for deployment in a wide range of devices across every industry segment.

Samsung NVMe Reference System

Samsung NVMe Reference System is a high-performance, all-flash NVMe scale-out storage server with up to 24 x 2.5" hot-pluggable Samsung advanced NVMe SSDs that provides extremely high capacity in a small footprint. It is based on PCIe Gen3 NVMe SSDs and offers the lowest latency in the industry with an optimized data path from the CPU to the SSDs.

Each SSD slot provides power and cooling for up to 25W per SSD to provide support of current and future generation large-capacity SSDs, as well as SSDs with different endurance and performance levels. With the PM953, the maximum capacity per system is 46TB. With the next-generation PM963 SSDs, the maximum capacity per system is 92TB. With the high-endurance PM1725a, the maximum capacity per system is 153TB. This is a dual-socket, Xeon®-based system and EIA-compliant 2RU chassis. It also uses 4x 40Gb/s networking connectivity with remote directory memory access (RDMA). The Samsung NVMe Reference System is available through StackVelocity (a business unit of Jabil Systems) as the Greyguard platform.

Having this balance means performance scales more linearly, without tending to be overprovisioned along any one of the components. With Ceph's distributed cluster capabilities, enterprises can now bring a performance tier reaching hundreds of thousands of IOPS to the traditional scale-out capacity tier that OpenStack offers.




Figure 2. Samsung NVMe Reference System

Features and capabilities

The features of the combined Red Hat Ceph Storage and Samsung NVMe Reference System and NVMe SSDs include:

- OpenStack integration
- S3 and Swift support using RESTful interfaces
 - High performance
 - 700K+ IOPS for small 4KB random IO across a three-node Ceph cluster
 - 30GB/s for large 128KB sequential IO across a three-node Ceph cluster
- Reference architecture based on the extensive testing jointly performed by Red Hat and Samsung to characterize optimized configuration
- Ability to use nonproprietary, commodity-based hardware
- Striping and replication across nodes to enable data durability, high availability, and high performance
- Automatic rebalancing using a peer-to-peer architecture to add instant capacity and data protection with minimal operational effort

- Upgrade clusters in phases – adding or replacing cards online – with minimal or no downtime
- Lower power consumption and higher reliability than similar capacity HDD configurations

28.5Gb/s for sequential read and 6.2Gb/s for sequential write

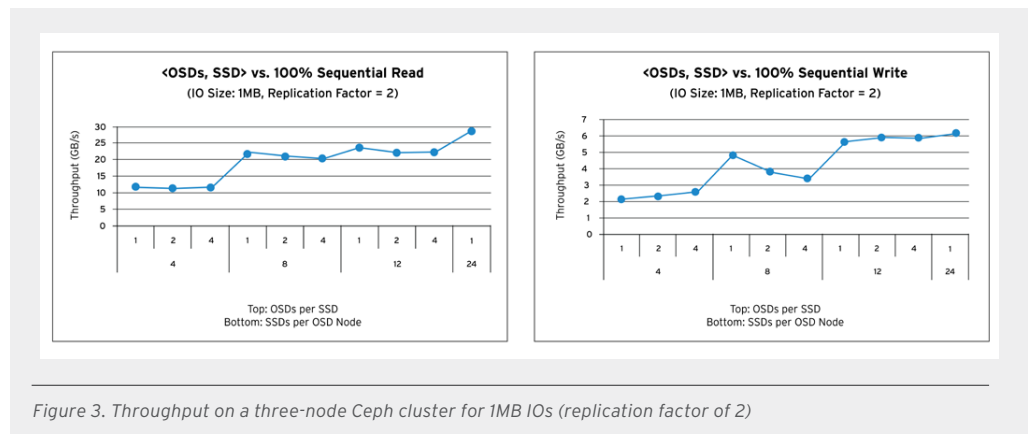


Figure 3. Throughput on a three-node Ceph cluster for 1MB IOs (replication factor of 2)

693K IOPS for random read and 87.8K IOPS for random write

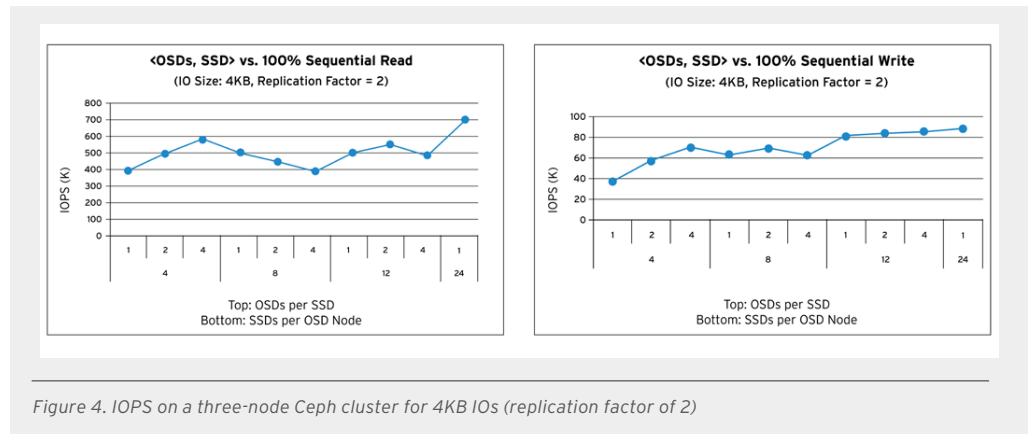
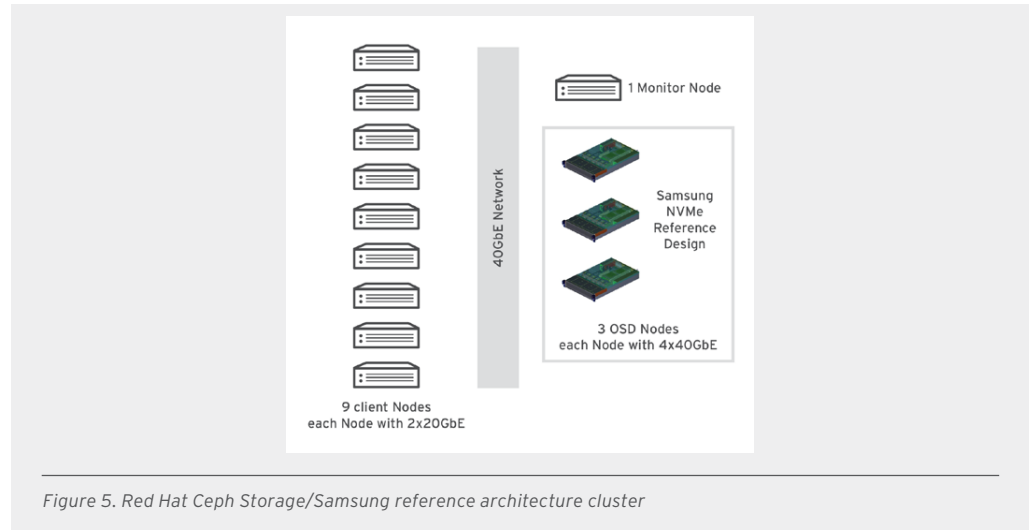


Figure 4. IOPS on a three-node Ceph cluster for 4KB IOs (replication factor of 2)

TECHNICAL DETAILS

Samsung and Red Hat have extensively tested a high-IOPS Red Hat Ceph Storage cluster running over the Samsung NVMe Reference System. This diagram shows the reference architecture configuration and the “speeds and feeds” metrics of the combined solution.



These tables provide the technical details of the combined solution. For more information about the Samsung NVMe Reference System, see Samsung’s All Flash Array reference design.¹

RED HAT CEPH STORAGE REFERENCE ARCHITECTURE

NVMe slots	<ul style="list-style-type: none"> • 24x 2.5" Samsung NVMe SSD slots • Each slot supports up to 25W per SSD • Support PM953 NVME SSD with maximum capacity of 1.92TB per SSD • Will support next-generation Samsung NVME SSDs, including PM963 and PM1725a
CPU	2x Intel [®] E5-2699 v3
Memory slots	Up to 512GB (minimum 256GB)
Network	Mellanox ConnectX-4 EN: 4X40GbE for network connectivity
Version of Red Hat Enterprise Linux and Ceph	Red Hat Enterprise Linux 7.2, Ceph Hammer LTS (0.94.5)
Number of Ceph nodes	3 all-flash NVMe storage nodes; 6 client nodes; 1 monitor node
Availability/redundancy	1+1 1200W power supplies, 4+1 redundant fans
Remote accessibility	Dedicated 1GbE BMC (KVM/IP, IPMI)
Form factor	2U EIA-310-D, L 28", H 3.43", W 17.15", UL, CE, FCC, RoHS

¹ <http://www.samsung.com/semiconductor/support/tools-utilities/All-Flash-Array-Reference-Design/>.

PERFORMANCE


IOPS (100% random read, IO size: 4KB)	696K
IOPS (100% random write, IO size: 4KB, replication factor: 2)	87.8K
IOPS (70% random read, 30% random write)	Read: 164.65K; write: 70.72K
Throughput (100% sequential read, IO size: 128KB)	30.74GB/s
Throughput (100% sequential read, IO size: 1MB, replication factor: 2)	6.76GB/s

CONCLUSION

Historically, enterprises have struggled to attain optimal OpenStack performance for certain use cases, such as distributed MySQL databases, telco NDVR long-tail content retrieval, and financial services. To overcome that challenge, Red Hat Storage and Samsung have partnered to create a solution that addresses several use cases that were previously out of the reach of traditional OpenStack deployments. A performance tier using Red Hat Ceph Storage and NVMe SSDs can now be deployed in OpenStack, supporting the bandwidth, latency, and IOPS requirements of high-performance workloads.

ABOUT SAMSUNG ELECTRONICS CO., LTD.

Samsung Electronics Co., Ltd. inspires the world and shapes the future with transformative ideas and technologies. The company is redefining the worlds of TVs, smartphones, wearable devices, tablets, cameras, digital appliances, printers, medical equipment, network systems, and semiconductor and LED solutions. For the latest news, please visit the Samsung Newsroom at news.samsung.com.



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

Copyright © 2016 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and JBoss are trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

The OpenStack® Word Mark and OpenStack Logo are either registered trademarks / service marks or trademarks / service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community.