



Look Inside.™

Better Together: Red Hat® Ceph Storage on Intel® Architecture

Delivering massive scalability and reduced costs for organizations adopting software-defined storage



Traditional storage can fall short of the mark for today's enterprise. As your organization moves into new, next-generation use cases such as cloud, analytics, big data, or media streaming, you need a storage solution that is economical to scale as data volumes grow. But traditional storage systems add fixed amounts of capacity with limited bandwidth, requiring a large, up-front capital expense and making scalability difficult.

With scale-out storage and standard, off-the-shelf servers, adding capacity and performance can be a smaller, incremental expense. Your organization can scale one server at a time. Cloud service providers and forward-looking enterprises are adopting this model, using open-source software designed to run on standard server hardware. Together, Intel and Red Hat provide the right storage combination, with the Red Hat® Ceph Storage software platform running on high-performing and reliable Intel® Xeon® processor-based servers with Intel® Ethernet Converged Network Adapters and the Intel® Solid-State Drives (Intel® SSDs) Data Center Family.



Ceph: Maximizing the Value of Software-Defined Storage

Red Hat Ceph Storage is a massively scalable, fully distributed, software-defined storage system developed from the ground up to deliver object and block storage in one self-managing, self-healing platform. Red Hat Ceph provides an excellent alternative to traditional storage appliances, one well-suited for analytics, big data, virtualization, and streaming media, and offers a unified solution for cloud computing environments. Ceph features such as erasure coding, cache tiering, and replication allow flexibility in how data is stored, protected, and accessed to meet varied use-case and performance needs.

Additional advantages of Ceph include no single point of failure, easy integration with third-party tools through application programming interfaces (APIs), and a peer-reviewed architecture that runs on standard server hardware. Figure 1 shows the elements of the Ceph storage architecture.

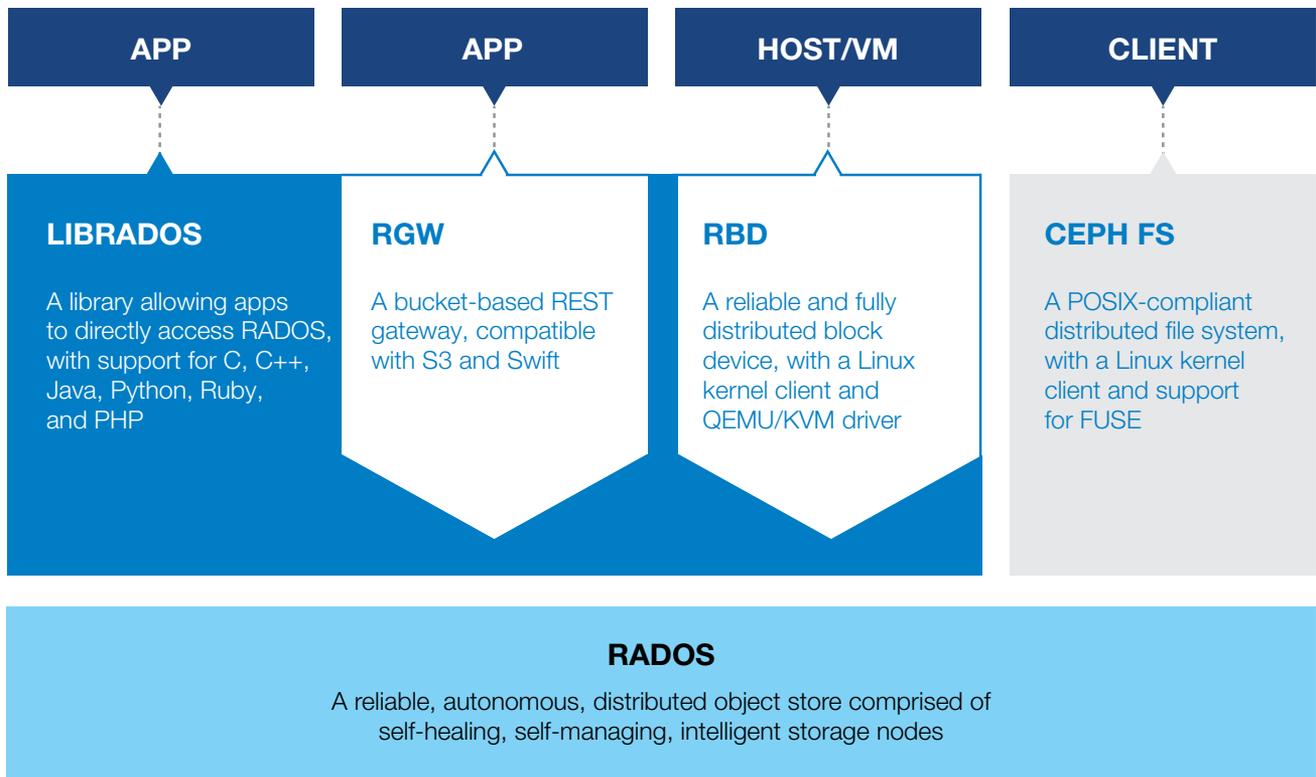


Figure 1. Elements of the Ceph storage architecture

Running on the Right Storage Hardware

With software-based storage, much of the value your organization obtains from a solution depends on the hardware you use to run it. For today's enterprises, it is important to use enterprise-ready, production-quality software on industry-standard hardware. Equally important is finding a provider that knows how to make reliable hardware that is well matched to the task. Intel and Red Hat have closely examined Ceph performance across varied Intel architecture-based hardware and used the capabilities of that hardware to make Red Hat Ceph Storage run even faster. Both Red Hat and Intel are active participants in the Ceph open-source community.

Harnessing Red Hat's Expertise and Support

Enterprises with successful software-defined storage deployments tend to partner with organizations that have deep open-source roots, tight alliances with community projects, and the ability to provide enterprise-level services and support. For software-defined storage, Red Hat is that company. Red Hat Ceph Storage offers the innovation of the Ceph community project while providing stability, expertise, and award-winning, enterprise-class support.

Combining Red Hat Ceph Storage with Intel Technology

As a result of their work together, Red Hat and Intel have developed a combined solution based on Ceph performance optimization and testing of Ceph with the following Intel technologies:

- **Intel Xeon E5 and E3 processor families:** These processor families provide energy-efficient high performance together with integrated Intel® Data Direct I/O technology to help remove bottlenecks, decrease latency, and increase data throughput. The processors also deliver significant benefits in scalability, reliability, and data protection. You benefit as well from Intel's regular cadence of processor performance and technology improvements.

- **Intel Solid-State Drive - Data Center Family:** Intel Data Center SSDs offer full end-to-end data protection, consistent performance with low latencies, high-write endurance, and scalability for growing storage needs. The SSDs provide a cache tier in a Ceph cluster and increase the speed on all writes when applied to the Ceph write journal. Intel has balanced fast read/write speeds with optimized CPU utilization for Ceph storage.
- **Intel Ethernet Converged Network Adapters:** Intel's unified networking vision is to support all LAN data and storage traffic on a common Ethernet infrastructure. 10 and 40 Gigabit Intel Ethernet Converged Network Adapters can help increase Ceph performance compared to traditional storage networking. Intel network adapters increase server bandwidth, providing scalable, high-throughput features while reducing power requirements.
- **Intel® Intelligent Storage Acceleration Library (Intel® ISA-L):** Intel ISA-L provides algorithms to help accelerate and optimize your storage on Intel architecture. The functions provided in this library help with data protection, data integrity, and data security. For example, using Intel ISA-L, Intel and the Ceph community accelerated erasure coding performance for Ceph on Intel systems by applying algorithms that speed functions through the CPU, thereby avoiding the typical performance trade-off of using erasure coding for data protection.

- **Virtual Storage Manager (VSM):** Intel has developed VSM open-source software to simplify the creation and management of Ceph clusters, and the open-source community continues to add refinements. VSM offers pre-designed cluster templates to make configuration easier. It also simplifies Ceph cluster management using a system to organize servers and storage devices according to performance characteristics and intended use. The VSM web-based user interface provides the operator with the ability to monitor overall cluster status, manage cluster hardware and storage capacity, and inspect detailed operation status of Ceph subsystems.

Accelerating Storage Performance

By moving to Ceph storage on Intel architecture-based servers, you benefit from the performance gains provided by the dedicated, expert resources of Red Hat and Intel. The two companies continually work together on tuning parameters and configurations, delivering improvements such as cache tier performance enhancements.

Reducing Operating Costs

Ceph is designed to be self-healing and self-managing, helping to reduce your administrator and budget overhead. Multiple components in the Ceph cluster can fail but the cluster will continue to function with the nodes dynamically rebalancing the data. In addition, the Controlled Replication Under Scalable Hashing (CRUSH) algorithm in Ceph frees storage clusters from the need for centralized data table mapping, eliminating this bandwidth-intensive task and single point of failure by distributing the mapping throughout the cluster.

The Ceph Community

Ceph is an open-source project built through the efforts of a dedicated, global community. The Ceph release series receives ongoing critical bug-fixes and stability improvements for a consistent and reliable deployment of Ceph.

Red Hat® Ceph Storage is based on this thriving Ceph open-source project. Both Red Hat and Intel participate actively in the Ceph community, contributing regularly to Ceph technology advancements and offering long-term commitment through patches and feature backports. Red Hat's proven track record in the open-source community and its award-winning support uniquely position it to represent customers' needs.

For more information about the Ceph community, visit: www.ceph.com

Scaling to Exabyte Level

CRUSH allows massive scalability for your organization's ever-growing volumes of storage. As nodes are added to the cluster, the expansion is automatically communicated to all parts of the system, and data is automatically redistributed across all nodes. CRUSH allows performance to scale linearly as a storage cluster grows by ensuring the client always retrieves data directly from the node it is stored on.

Helping Ensure Data Protection

With Ceph, your data is protected in multiple ways. Data placement through topology-aware CRUSH maps helps minimize failure-domain risks. Replication allows your organization to protect data against the failure of any one drive by placing replicated copies on different drives in the cluster. Alternatively, Ceph erasure coding enables data to be equally protected, but more efficiently.

Moving Beyond Storage Limitations

While much of today's storage is still based on expensive, closed infrastructure siloed from other storage resources, servers have moved beyond those limitations. It's time for storage to catch up. That's why forward-looking organizations are moving to software-defined, server-based storage solutions. The shared commitment of Red Hat and Intel to advance Ceph and optimize it on Intel architecture can help your organization make the move successfully.

Red Hat Ceph Storage enhances the economics of storage and provides a foundation for managing exponential data growth. It combines the most stable version of Ceph for object and block storage with the hardware and tools you need to run an enterprise production Ceph cluster at scale with confidence. Let Red Hat Ceph Storage on Intel architecture transform your organization's ability to manage today's vast and growing amounts of data.

For More Information

To learn more about Red Hat Ceph Storage, visit: www.redhat.com/en/technologies/storage/ceph

For more information about Intel technologies, visit: www.intel.com/storage



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel® products are not intended for use in medical, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

© 2015 Intel Corporation. All rights reserved. Intel, the Intel logo, Look Inside, the Look Inside logo, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.

Printed in USA

0415/CM/TDA/XX/PDF

Please Recycle

332147-001US