




Storage Wars: SDS and the inevitability of Ceph

Why Ceph will emerge as the victor in the SDS revolution.

By Phil Straw
CTO SoftIron



Software-defined storage (SDS) is the way of the future for enterprise – and Ceph will dominate SDS.

Last year, Gartner research revealed that SDS solutions “can reduce TCO by 50% or more without sacrificing performance, data services software robustness, or availability service-level objectives” compared with branded integrated legacy external controller-based (ECB) storage offerings.¹ Fifty percent! Further, “when deployed as the storage element of a hyper-converged integrated system reference architecture, cost savings can approach 70% compared with a traditional three-tier storage-area-network-based legacy ECB storage infrastructure.” In regard to a strategic planning, the analysis shows that “by 2019, 50% of existing storage array products will also be available as software-only versions...[and] by 2020, 70% of storage provisioning and management functions will be integrated into the infrastructure platform,” up from 10% just last year.² Software really is eating the world, and while storage comes with unique considerations, it is not exempt.

SDS as a concept — policy-based data storage provisioning and management independent of underlying hardware — is driven not only by enterprises’ eternal need to reduce costs, but also by the desire to enable the wide range of emerging capabilities like distributed, scale-out storage for easy growth and automatic redundancy and replication for data protection, as well as the natural repulsion toward vendor lock-in.

SDS is the “New Normal”

The capital cost requirements for purchasing a high-end enterprise storage system from Pure, Nimble, HPE, Dell/EMC, etc. is enormous — usually exceeding \$100k just to get started in equipment costs and long-term support agreements. Meanwhile, with SDS, you can run a multi-featured, highly configurable, and less-expensive storage platform on equally less-expensive hardware. Software-defined storage has started to get to a point where, instead of RAID hardware or silicon assets or high speed feeds, the idea is you aggregate horizontally — many, many components — and the more components, the more improvement. It is faster. It is more resilient. As you grow horizontally, you provide abstractions away from the hardware and the aggregation, and you achieve all the features and the performance that you could get from any traditional

solution. It’s not that the traditional sources aren’t making good storage solutions, it’s just that those solutions are outrageously exorbitant and monolithic compared to the accessible and flexible “new normal” presented by SDS.

This sea change was triggered by the rising tide of the agile, adaptable, continuous-improvement-charged ethos of the open source movement and, as cloud and storage consultant Jim O’Reilly noted in Network Computing last fall, “We are about to see a substantial portion of the IT industry migrating to the open source approach. While this isn’t good news for the traditional vendors, I expect we’ll see something like the Linux revolution taking place in storage, with the result a much healthier industry driven by a high rate of innovation.”³

Ceph will be the next Linux

SDS is the future, and that reference to the Linux revolution explains why Ceph is going to sweep this storage contest.

Back in 1991, when University of Helsinki computer science student Linus Torvalds posted about his OS kernel “hobby,” no one predicted that it would become what is arguably the world’s most ubiquitous bit of software. And this ascendance came to pass in the face of persistent opposition by seemingly invincible and resource-rich corporate behemoths.⁴

Without delving into every contributing factor to the Linux success story, it’s safe to say it “won” because it is open, flexible, and widely embraced by the global developer community. And herein lies the argument for why Ceph will fulfill a Linux-like destiny in open-source software-defined storage, because it delivers a similarly feature rich and extremely flexible platform that is already widely embraced.

Linux “won” because it’s open, flexible, and widely embraced by the global developer community. Ceph fulfills a Linux-like destiny in open-source software-defined storage because it delivers a similarly feature rich and extremely flexible platform that is already widely embraced.

Like Linux, Ceph’s origin traces back to a university student, Sage Weil at U.C. Santa Cruz working on his doctoral thesis. It began a decade ago as a pure mathematics technique that allows distribution, error recovery, and encoding such that you get more storage bang for your buck, with resilience thrown in to boot. Basically, the idea was that you could store more data (via pseudo-compression) and also get duplication, independent of underlying hardware. As Linux Journal summarizes, “In the world of Ceph, data is treated and stored like objects. This is unlike traditional (and legacy) data storage solutions,

¹ “High-Tech Tuesday Webinar: Making the Case for Software-Defined Storage in 2016”, Gartner, June 7, 2016.

² “Gartner Report on Software-Defined Storage”, Storage Newsletter, September 9, 2016.

³ “Best Open Source Storage Tools for the Enterprise”, Network Computing, August 29, 2016.

where data is written to and read from the storage volumes via sectors and at sector offsets (often referred to as blocks). When dealing with large amounts of large data, treating them as objects is the way to do it. It is also much easier to manage. In fact, this is how the cloud functions — with objects. This object-driven model enables simplified scalability to meet consumer demand easily. These objects are replicated across an entire cluster of nodes, giving Ceph its fault-tolerance and further reducing single points of failure.”⁵

SDS and Ceph: Optimized for Content Delivery, Archival Storage, Block Storage for Virtualization

Ceph, as a distributed, scale-out storage framework, caters best to high bandwidth, medium latency types of applications, such as content delivery (think Netflix, Comcast, AT&T), archive storage (Dropbox-type applications), or block storage for virtualization, but it can handle almost anything. In addition, Ceph can do all three types of commonly found storage types (object, block, and file) — a flexibility that separates it from the rest of the SDS herd. It’s inherent scale-out support aligns well with horizontal which means you can gradually build large systems as cost and demand require, and it sports enterprise-grade features such as erasure coding, thin provisioning, cloning, load-balancing, automated tiering between flash and hard drives, and simplified maintenance and debugging. Its Network File System (NFS) interface allows access to the same data from both Amazon S3 and NFS file interfaces and is compatible with the Hadoop S3A filesystem client, enabling developers to use Apache Hadoop MapReduce, Hive, and Spark. It is the Swiss Army Knife of the SDS world.


Additionally, Ceph’s reputation and user base have steadily grown as the SDS ecosystem continues to mature. Development has been shepherded by Red Hat since its 2014 acquisition of Weil’s startup, and the Ceph Community Advisory Board was founded in 2015 to seize the leadership mantle for open source software-defined storage technology. The charter advisory board includes rival companies cooperating as Ceph enthusiasts, such as Red Hat,

SUSE, and Canonical (Ubuntu Linux’s parent company), as well as impressive boosterism from CERN, Cisco, Fujitsu, Intel, and SanDisk. In further evidence of Ceph’s expanding reach, a recent OpenStack user survey shows Ceph dominating deployment — accounting for more than half of the block storage drivers in use.⁶

There are plenty of other SDS frameworks vying for mindshare, both proprietary and open source, each with their own strengths and weaknesses — and more will undoubtedly emerge as the shift away from traditional enterprise storage models intensifies. But Ceph has that Linux-like magic combination of qualities that heralds its destiny: accessibility, openness, unparalleled flexibility, rich features, an impressive and growing roster of enterprise ambassadors, and an already deep and devoted user community. Like Linux, Ceph is inevitable.

SoftIron® makes the world’s finest solutions for the data center. The company’s HyperDrive™ software defined storage portfolio is built on Ceph and runs at wire speed, while HyperCast™ delivers the best density and value for real-time video streaming. SoftIron unlocks greater business value for enterprises by delivering great products without software and hardware lock-in.

Copyright © SoftIron Limited, 2017. All rights reserved. SoftIron, HyperDrive, HyperCast and the SoftIron logo are registered trademarks of SoftIron Limited. ARM is a registered trademark of ARM Limited (or its subsidiaries) in the EU and/or elsewhere. AMD, the AMD arrow logo, and combinations thereof are trademarks of Advanced Micro Devices Inc. Socionext is a registered trademark of Socionext, Inc. SoftIron disclaims proprietary interest in the marks and names of others. This document is for information only. No warranties are given or implied. Contents are subject to change without notice. SoftIron Limited is registered in England at One Mayfair Place, London W1J 8AJ United Kingdom.



⁴ "MS' Ballmer: Linux is Communism", The Register, July 31, 2000.

⁵ "Understanding Ceph and Its Place in the Market", Linux Journal, June 21, 2016.

⁶ "OpenStack User Survey", April 2016.