



Storage Wars: Ceph powers the most exciting research on Earth

How researchers around the world are using Ceph at scale

By Harry Richardson
Chief Scientist, SoftIron



As researchers seek scalable, high performance methods for storing data, Ceph is a powerful technology that needs to be at the top of their list.

Ceph is the leading open source software-defined storage platform and while it's not often in the spotlight, it's working hard behind the scenes, playing a crucial role in enabling ambitious, world renowned projects such as CERN's particle physics research, ImmunityBio's cancer research, The Human Brain Project, The MeerKAT radio telescope, and more.

These ventures are propelling the collective understanding of our planet and the human race beyond imaginable realms and the outcomes will forever change how we perceive our existence and potential. It's high-time Ceph receives the praise it deserves for powering some of the most exciting research projects on Earth.

Why Ceph for Research?

Ceph is flexible, inexpensive, fault tolerant, hardware neutral, and infinitely scalable which makes it an excellent choice for research institutions of any size. "Ceph has the capability to support research at any level," says Phil Straw, CEO at SoffIron. "Many research organizations have unique, complex storage requirements and don't want to be locked in to a particular hardware vendor. Ceph is a great fit for them."

Ceph's benefits for researchers include:

- » **Supports multiple storage types including object, block, and file systems** – Regardless of the type of research being conducted, the resulting files, blocks and/or objects can all live in harmony in Ceph.
- » **Hybrid cloud-ready** – Ceph natively supports hybrid cloud environments which makes it easy for remote researchers, who might be located anywhere in the world, to upload their data in different storage formats.
- » **Hardware-neutral** – Ceph doesn't require highly performant hardware, which lowers equipment costs and eliminates vendor lock-in.
- » **Resilient** – There's no need to buy redundant hardware in case a component fails, because Ceph's self-healing functionality quickly replicates the failed node, ensuring data redundancy and higher availability.

In this article, we'll examine how four organizations with vastly different research projects and unique data storage requirements are using Ceph.

CERN

Scientists from around the globe use CERN's particle accelerators to explore questions such as "What is the nature of our universe?". CERN's super-sized data center executes more than 500,000 physics jobs daily¹ and its current storage requirements are estimated to be 70 petabytes per year.² CERN selected Ceph because of its ability to build block storage for OpenStack, and the fact that remote servers can easily be added with no downtime.³

ImmunityBio

Genomics research requires the manipulation of massive amounts of data. ImmunityBio, a leader in molecular testing and personalized cancer treatments, processes enormous amounts of data, including one terabyte per genetic test, so it's important that storage should not become a bottleneck. It takes one month to process raw data on an 800-core cluster, and the workload can vary from 2.5 million small random files to a handful of giant, sequential files. To make its storage requirements even more complex, ImmunityBio's data is "infinitely useful", meaning it will be stored forever for future research or reprocessing.

ImmunityBio chose Ceph because it is very good at processing and storing large amounts of data cost-effectively. The fact that Ceph supports unified storage of object, block and file types, and NantOmics can manage Ceph itself without relying on an outside vendor were also attractive.

Even though the cloud is a popular choice for storage, ImmunityBio avoided that option, because it believes cloud pricing isn't scalable. Cloud vendor lock-in is also an issue, because it's notoriously difficult to move one petabyte of data between cloud vendors.

With Ceph, ImmunityBio has achieved cost-effective storage, better performance and reliability, and eliminated vendor lock-in, allowing them to pursue their research unhindered.

1. OpenStack at CERN: A 5 Year Perspective, 2018.

2. CERN: The secret sauce, Data Center Dynamics, October 25, 2018.

3. At CERN, storage is the key to the universe, Superuser, June 22, 2015.



Human Brain Project

Re-creating the intricate, complex processes of the human brain using technology is, by any definition, a massive undertaking.

The [Human Brain Project](#) (HBP) is a ten-year European Union flagship research project, based on exascale supercomputers, that aims to advance knowledge in neuroscience, computing and brain-related medicine.⁴

One of the goals of the HBP is to provide researchers worldwide with tools and mathematical models for sharing and analyzing data to understand how the brain works in order to emulate its computational capabilities.⁵ The scale of this project is hard to fathom: the human brain is so complex that a normal computer can't simulate even a fraction of it – in fact, just one of the supercomputers used in the HBP is as powerful as 350,000 standard computers.

A significant portion of the HBP uses massively parallel applications in neuro-simulation to interpret data. The advanced requirements of the HBP are far beyond current technological capabilities and will undoubtedly drive innovation in the high performance computing industry.

The HBP utilizes a next-generation storage system based on [Ceph](#), exploiting complex memory hierarchies and enabling next-generation mixed workload execution.⁶ With Ceph, the HBP eliminates vendor lock-in while realizing 90% read efficiency and demonstrating outstanding scalability as object sizes increase.

MeerKAT Radio Telescope

Imagine an array of telescopes, collecting vast amounts of information about outer space, located in the world's most remote and harshest locations. The [MeerKAT radio telescope](#) is a 64-antenna array radio telescope, built on the Square Kilometre Array (SKA) site. The SKA project is an international effort to develop the world's largest radio telescope with a square kilometer of collecting area. Ceph is used to store and retrieve huge volumes of data including a 20-petabyte object-based storage system.

One of MeerKAT's unique challenges is the isolated location of the telescope arrays: one is located deep in the South African desert and the other in the Australian outback. Cost is a key factor, because the project requires a massive amount of storage hardware. Ceph is an excellent choice, because it doesn't require highly performant, expensive hardware for optimal performance.

Resiliency is also critical, because the telescopes are located in remote environments which makes it difficult to quickly procure new hardware if a component fails. If a node failure does occur, Ceph's self-healing functionality quickly replicates the failed node using secondary copies located on other nodes in the cluster, thereby ensuring data redundancy and higher data availability. As a result, MeerKAT has a highly resilient, scalable storage solution that maximizes efficiency while minimizing costs.

Replicating data from each of MeerKAT's far-flung locations to a centralized data store is also critical. Using Ceph, the data for each telescope array is replicated to a centralized data store in Cambridge, England that's part of the Square Kilometre Array project. This allows all the MeerKAT data to be analyzed in its entirety while ensuring availability.

Ceph's Inherent Hardware-Neutrality

One of the common issues for each of the research projects highlighted in this article is storage hardware. Ceph's inherent hardware neutrality is a great benefit to researchers as they aren't limited to proprietary hardware solutions that are often expensive and inflexible.

Ceph is so versatile that it can be run on nearly anything: a server, a Raspberry Pi, even a toaster (assuming it runs Linux). For research purposes, scientists can choose to run Ceph on a black box server, or they could use HyperDrive, a purpose-built storage appliance, built by SoftIron, that's optimized specifically for Ceph. Research institutions like the University of Minnesota's Supercomputing Institute and ImmunityBio, are realizing the added benefits of using a custom-designed, optimized Ceph storage appliance like HyperDrive, to power some of the most exciting research projects on Earth.

To learn more about SoftIron, Ceph or HyperDrive, visit softiron.com

4. Human Brain Project, Framework Partnership Agreement, 2015.

5. Ceph on the Brain: A Year with the Human Brain Project, StackHPC Ltd.


6. Ceph on the Brain: Storage and Data-Movement Supporting the Human Brain Project, insideHPC, April 17, 2018.



SoftIron® makes the world's finest solutions for the data center.

The company's **HyperDrive®** software-defined storage portfolio is built on Ceph; it's custom-designed and purpose-built for scale-out enterprise storage and runs at wire speed.

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.

 softiron.com

 @SoftIron

 @SoftIron

 @SoftIronNews

 info@softiron.com

Copyright © SoftIron Limited, 2019. 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.