



NetApp®

Success Story

Poland's National Centre for Nuclear Research Accelerates HPC with NetApp Clustered Data ONTAP



KEY HIGHLIGHTS

Industry
Energy

The Challenge

Meet speed and reliability requirements for file-based scientific workloads with a single, unified storage platform to achieve research results faster.

The Solution

Migrate data to a two-node NetApp® FAS6280 storage system running NetApp clustered Data ONTAP® and use NetApp Flash Cache™ to accelerate performance.

Benefits

- Process experiments 24/7 with no downtime, accelerating research
- Streamline workflows for scientists
- Save US \$10,000 per downtime incident avoided, enabling more efficient use of funds
- Meet throughput requirements in a 40% smaller storage footprint

Customer Profile

The largest research institute in Poland, the National Centre for Nuclear Research (Narodowe Centrum Badan Jadrowych, or NCBJ) was established in 2011 to develop nuclear power technologies and promote practical and safe applications of nuclear physics. NCBJ conducts research to support decision making about how to best develop an economical and environmentally sound nuclear power industry in Poland.

With more than 1,000 employees, NCBJ operates the MARIA nuclear reactor and also collaborates with international research institutes, including CERN in Geneva. The NCBJ Swierk Computing Centre provides the high-performance computing (HPC) infrastructure for much of the institute's research.

The Challenge

Providing researchers with reliable, fast storage

In addition to energy research and particle physics, NCBJ monitors and simulates the spread of radiation, chemicals, and pollution by analyzing data collected from sensors. All of these research activities require fast, reliable storage—and lots of it. The

institute has accumulated more than 1PB of unstructured file data, and after the data is stored, it must be quickly processed so that researchers can draw conclusions and carry on with their experiments.

“We need capacity, performance, and high availability from our storage,” says Adam Padee, head of computational infrastructure at NCBJ. “Any downtime is very expensive for us, and frustrating for our scientists. In addition, any infrastructure failure necessitates that all running compute jobs be restarted from the beginning; in some cases, this can mean losing weeks’ worth of compute time.”

Until recently, occasional storage failures were a reality. A legacy storage array running the Lustre file system provided speed for file-based scientific workloads, but it could not reliably handle the demands of NCBJ’s HPC environment. As soon as a compute job was completed, researchers had to move the data to more reliable storage.

To get around this challenge, many researchers chose to run their compute jobs on laptops or from their home directories hosted on HP storage. “The HP storage was reliable, but it was too

slow for our needs,” says Padee. “Our researchers ended up trading performance for reliability, and we didn’t want to put them in the position where they had to make that choice.”

The Solution

Nondisruptive operations to support critical research

With assistance from NetApp and NetApp Partner Koma Nord, NCBJ deployed a two-node NetApp FAS6280 storage system running the clustered Data ONTAP operating system to eliminate planned and unplanned downtime. HP servers running Scientific Linux® connect to the NetApp cluster using the NFS protocol over a Juniper Networks network infrastructure.

NetApp Global Services performed the clustered Data ONTAP Migration Service and moved users’ home directories to the NetApp cluster, and Koma Nord provided systems integration assistance. “Koma Nord was extremely helpful and cooperated well with the NetApp Global Services team,” says Padee.

Having a unified cluster architecture such as clustered Data ONTAP is particularly valuable to organizations like NCBJ that must maintain constant availability for large amounts of scientific data. All of the NetApp storage systems can be managed as a single logical pool that can seamlessly scale to tens of petabytes and thousands of volumes, with a global namespace for easy workload mobility.

Both storage controllers are equipped with NetApp Flash Cache PCIe-based intelligent caching, which helps optimize storage performance, improve storage efficiency, and reduce costs. With Flash Cache, NCBJ has the potential to increase I/O throughput by up to 75% for random read-intensive workloads and can reduce latency by a factor of 10 or greater.

“Achieving reliability and speed with a single storage platform has traditionally been challenging for us,” says Padee. “The combination of NetApp clustered Data ONTAP and Flash Cache solved that problem.”

NCBJ uses NetApp Snapshot™ and SnapRestore® technologies to back up the contents of users’ home directories, mitigating the risk of data loss, and NetApp OnCommand® Unified Manager simplifies storage management.

Business Benefits

Streamlining scientific workflows

With clustered Data ONTAP, scientists no longer need to compromise performance for reliability or worry about where data should reside. They simply run all compute jobs from their home directories on the NetApp cluster, confident that even if one storage controller goes offline, jobs will continue to run. Scheduled downtime windows are no longer required.

“Our researchers have a lot more confidence in our HPC infrastructure now that we’ve moved to NetApp clustered Data ONTAP,” says Padee. “We’re saving money as well—every time we previously had a failure, it cost us at least US \$10,000 in lost productivity, compute time, and resources.”

Engineers noticed a significant boost in performance with NetApp Flash Cache. “Using NetApp Flash Cache allowed us to meet our throughput requirements with a 40% smaller storage footprint,” says Padee. “Flash Cache is particularly effective with computational fluid dynamics applications—in many cases, we would need threefold more disk to achieve equivalent performance without Flash Cache.”

Speeding time to discovery

Allowing researchers to run their compute jobs on high-performance storage without worrying about interruptions or

data loss is helping NCBJ achieve its objectives as a world-class research organization.

“NetApp Clustered Data ONTAP is helping us accelerate the pace of our research to discover new ways that Poland can benefit from nuclear power and other energy sources,” says Padee. “We now have performance, stability, and simplicity in a single system, which makes life a lot easier for both IT and our scientists.”

SOLUTION COMPONENTS

NetApp Products

FAS6280 storage system

Clustered Data ONTAP 8.2P3

OnCommand Unified Manager 5.1

Flash Cache

Snapshot and SnapRestore technologies

Environment

Applications: Open-source and internally developed job scheduling and software configuration management tools; ANSYS engineering simulation software; OpenFOAM CFD; RELAP transient analysis software; GEANT detector description and simulation tool; FLUKA particle physics simulation package

Server platform: HP servers running Scientific Linux

Network: Juniper Networks

Protocol
NFS

NetApp Global Services

Clustered Data ONTAP Migration Service

Partner

Koma Nord Sp. z o.o.
www.komanord.pl



www.netapp.com

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Discover our passion for helping companies around the world go further, faster at www.netapp.com.

Go further, faster®

© 2014 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, Flash Cache, OnCommand, SnapRestore, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. CSS-6690-0214

Follow us on: