Quantum

Image courtesy of NASA

CASE STUDY

Brrrrrrr: CReSIS Measures the Impact of Sea Level Change on Ice Sheets with Quantum StorNext

Led by a team at the University of Kansas, the Center for Remote Sensing of Ice Sheets (CReSIS) uses advanced radar devices to image and track changes in ice sheets in Greenland and Antarctica. The organization chose Quantum StorNext scale-out storage to support tremendous volumes of scientific data on its HPC cluster.

CReSIS
Center for Remote Sensing of Ice Sheets

FEATURED PRODUCTS

StorNext 5

> With the StorNext platform, we now have a single file system that manages our entire multi-tier storage solution, from tape to primary disk to the HPC cluster.

Riley Epperson
IT Engineer, Center for Remote Sensing of Ice Sheets, University of Kansas

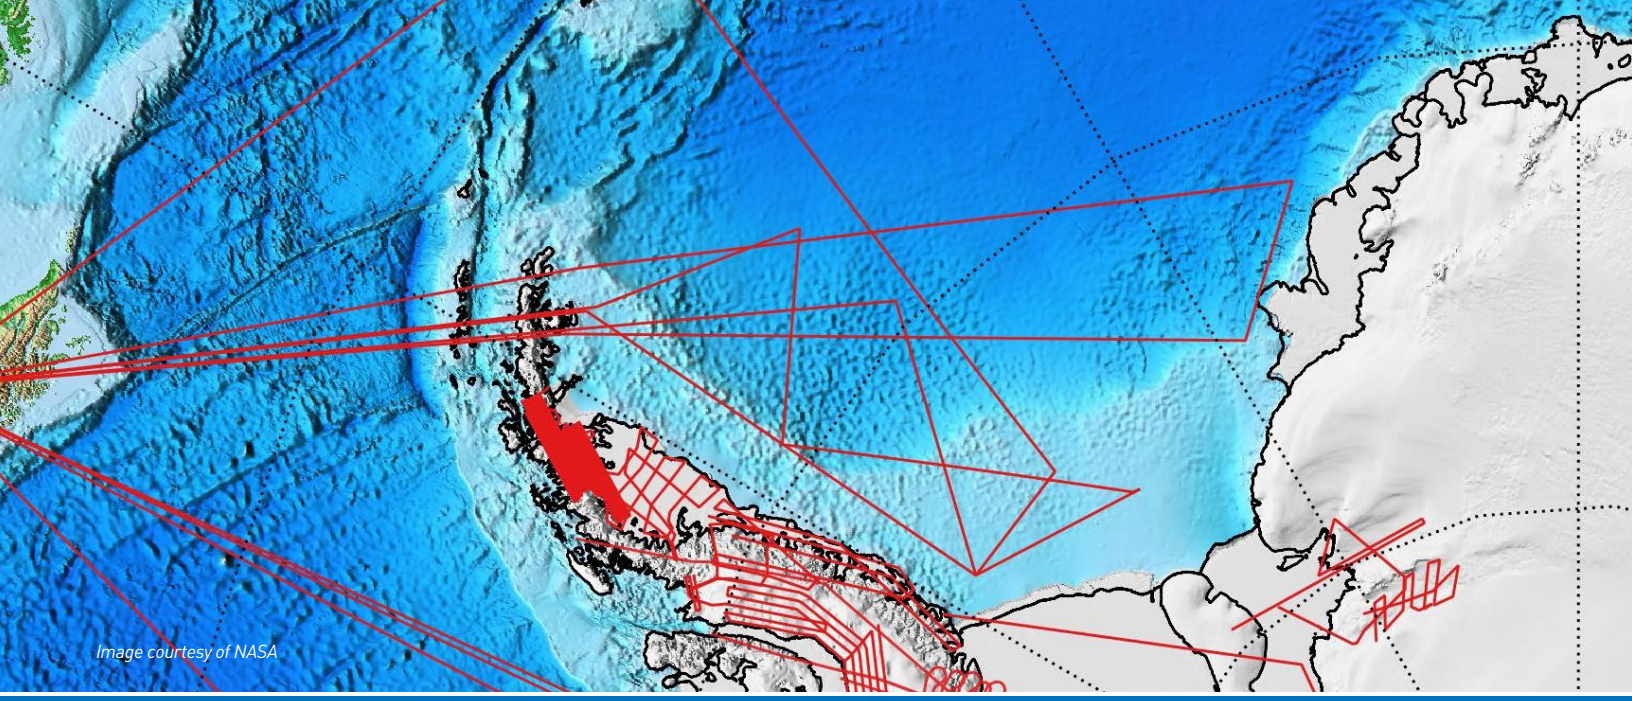

Image courtesy of NASA

> We have more than tripled our storage performance compared with the previous environment.... StorNext helps us ingest and process growing collections of data faster while continuing to meet our existing deadlines.

Riley Epperson - IT Engineer, CReSIS, University of Kansas

CReSIS
Center for Remote Sensing of Ice Sheets

SOLUTION OVERVIEW

- Quantum StorNext® Scale-out Storage
- StorNext Metadata Appliance
- Primary disk storage
- Tape archives

KEY BENEFITS

- **Scales to support increasing data growth** as scientists collect more ice data per mission
- **Triples storage performance** to speed workflows even as data grows
- **Saves on costs as storage expands** by enabling greater use of economical tape
- **Provides a multi-tier storage environment** without adding administrative complexity

Established in 2005 by an award from the National Science Foundation, CReSIS helps precisely measure changes in the earth's ice sheets. "By mapping the layers, the bed, and the top of glaciers in Greenland and Antarctica, scientists can assess current effects of climate change on ice and predict its future impact," says Riley Epperson, IT engineer for CReSIS.

CReSIS is led by an interdisciplinary group of engineers, geologists, glaciologists, and other scientists from the University of Kansas (KU). Those researchers regularly collaborate with university researchers and government agencies around the world. "We work with scientists from research institutions such as the University of Washington, the University of Copenhagen, and Naval Research Laboratories," says Epperson. "With NASA, we participate in Operation IceBridge—the biggest non-satellite project to map ice in Antarctica and Greenland."

To examine the ice sheets, CReSIS uses advanced radar, often attached to the underbellies of planes. "For Operation IceBridge, researchers used radar systems installed on DC-8 aircraft," says Epperson. "Research teams would take off from Chile, collect data over Antarctica for eight hours, then fly home. They would repeat that process five days a week, weather permitting, until the mission was complete."

COLLECTING UP TO 150TB OF SCIENTIFIC DATA PER MISSION

The amount of data collected per mission has grown dramatically over time. "In 2005, researchers might have collected 100GB per mission," says Epperson. "But in the years since then, our engineers have continued to enhance the capabilities of radar systems. Today's systems can provide much greater resolution of the ice layers. As a result, researchers now bring home closer to 150TB per mission."

Theoretically, researchers could collect even more data. "It's more of a financial limit than anything else," says Epperson. "Researchers use hard disk drives and solid-state drives to

collect data in the field. We purchase as many disks for researchers as the budget allows."

As the amounts of raw data collected in the field grow, it can hinder the processing of the data for the scientists who run analyses and create the final research products. "The more data researchers collect, the more time we spend ingesting, verifying, and cleansing that data," says Epperson. "We needed a solution that would help us improve the efficiency of our workflow."

SEARCHING FOR A HIGH-PERFORMANCE STORAGE PLATFORM THAT SCALES

A few years ago, the CReSIS IT group decided to change the file system it was using to store and manage research data. "The hardware vendor for our disk-based storage system was ending support for its file system," says Epperson. "We needed a new approach to the large and growing data volumes while also allowing us to continue to use our existing storage."

In choosing a new file system, the IT group wanted a solution that could help increase throughput. "The previous file system limited the storage pool size to 256TB. That restriction limited the number of disks we could place in each pool, and as a result, it limited our IOPS," says Epperson. "We wanted a file system that would help us increase throughput by allowing us to use more drives for each storage pool."

High throughput is important for meeting rigorous service-level agreements. "Collaborators and partners need to analyze and process the data that we ingest so they can generate the products required for research," says Epperson. "With some partners, such as NASA, we are expected to turn around data within a particular timeframe."

REQUIRING A SOLUTION THAT SCALES MORE COST-EFFECTIVELY

The CReSIS IT group also needed data management capabilities that could support a multi-tier environment. "We ingest raw data only two to three times per year, then it just sits—approximately 80 percent of our data is cold. There is no point in keeping that data on expensive disk if it's not being frequently accessed," says Epperson. "We wanted a solution that would enable us to easily move cold data to tape and then retrieve it later on."

Simplifying those processes was essential for controlling administrative burdens. "We have a very small team," says Epperson. "We need data management that simplifies and streamlines tasks so we don't have to spend too much time on things like data movement."

BUILDING A MORE FLEXIBLE, LARGE-SCALE STORAGE ENVIRONMENT WITH STORNEXT

After evaluating a variety of new file systems and data management solutions, the CReSIS team selected the Quantum StorNext platform. "We reached out to several other Quantum StorNext customers, and heard very positive feedback," says Epperson.

The team implemented a StorNext multi-tier solution to better manage the growing collection of scientific data gathered from CReSIS missions. Powered by StorNext data management software, the solution includes a high-performance shared file system and policy-driven tiering software.

Selecting the StorNext platform enabled the CReSIS team to continue to use their aging primary disk. "We did not have to rip and replace anything—we could protect our existing investment," says Epperson. "At the same time, we could easily add a tape archive tier into the storage solution. With the StorNext platform, we now have a single file system that manages our entire multi-tier storage solution, from tape to primary disk to the HPC cluster."

SPEEDING UP INGEST AND ENABLING HPC WORKFLOWS FOR CLIMATE RESEARCH

Once a mission is complete, researchers deliver disk drives to the CReSIS IT team. The team then ingests the data into its HPC cluster to cleanse the data, filtering out any noise from the raw data. "StorNext eliminates the pool size limitations imposed by our previous file system," says Epperson. "We currently have 600 disks on a single file system instead of only 100."

> "Each time engineers create a new radar system, we might need to store five times as much data per mission. With StorNext we can quickly scale to meet those requirements."

Riley Epperson,
IT Engineer,
CReSIS

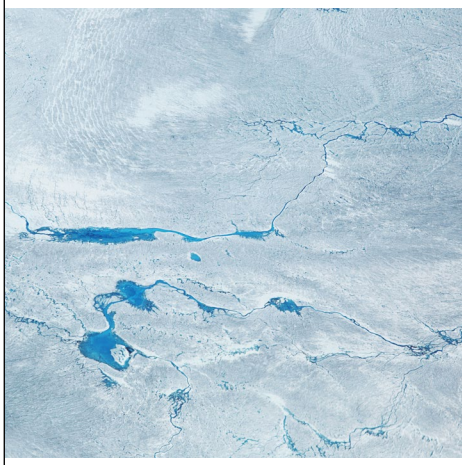

ABOUT CRESIS

The Center for Remote Sensing of Ice Sheets (CReSIS) is a Science and Technology Center that was established by the National Science Foundation in 2005. The organization strives to make meaningful contributions to the ongoing, urgent work of addressing climate change. It develops new technologies and computer models that measure and predict the response of sea level change to the mass balance of ice sheets in Greenland and Antarctica. To support new scientific research, the CReSIS IT group stores, manages, and provides access to large volumes of data collected through ice sheet exploration.

The result is superior performance. "We have more than tripled our storage performance compared with the previous environment," says Epperson. "StorNext helps us ingest and process growing collections of data faster while continuing to meet our existing deadlines."

High-speed storage helps speed up the overall workflow. "The faster we can move data into our HPC cluster for processing, the faster we can deliver it to researchers," says Epperson.

CONTROLLING COSTS WITH A MULTI-TIER STORAGE SOLUTION

Moving to the StorNext platform has enabled CReSIS to expand its storage environment while keeping costs under control. "We were maxing out our disk storage and didn't want to buy more cabinets—especially since so much of the data we collect stays cold," says Epperson. "StorNext allows us to scale our environment using economical tape archives and still manage all our data within a single namespace."

The cost savings from implementing a multi-tier solution that includes tape as well as disk has added up quickly. "StorNext offers a strong return on investment," says Epperson.

StorNext also provides the flexibility to meet unpredictable changes in storage requirements. "We don't have a consistent rate of data growth from year to year," says Epperson. "Each time engineers create a new radar system, we might need to store five times as much data per mission. With StorNext we can quickly scale to meet those requirements."

ENVISIONING ADDITIONAL USES FOR STORNEXT AT UNIVERSITY OF KANSAS

Given the positive benefits that CReSIS has experienced so far with the StorNext platform, it's not surprising that the team recommends StorNext to others. "Organizations running HPC environments might not immediately think of StorNext, but they should," says Epperson. "The StorNext platform can help deliver the speed that organizations need to accelerate research."

The IT group also sees a bright future for StorNext across KU. "Administrators began looking for a campus-wide storage platform that could support a broad range of research projects: We recommended StorNext," says Epperson. "It's a reliable platform that provides a flexible, shared file system and tiering capabilities that could work well for supporting multiple departments at KU and elsewhere."

To contact your local sales office, please visit www.quantum.com

Quantum

CS00377A-v01