University of Groningen breaks new ground in scientific research

With high-performance technology from IBM

university of groningen

Overview

The need

To maintain its position at the forefront of cutting-edge research, the University of Groningen wanted to develop a high-performance computing solution to support large-scale research projects.

The solution

The University of Groningen partnered with IBM® Global Technology Services® to build the Target platform: a state-of-the-art infrastructure for processing and storing vast amounts of data. The solution leverages IBM RackSwitch™ technology for optimal network performance.

The benefit

Raises network performance, boosting bandwidth levels by 50 percent; supports ground-breaking research initiatives; provides high levels of scalability to accommodate rapid growth in data volumes.

Founded in 1614, the University of Groningen in the Netherlands is one of the oldest and largest research universities in Europe. In 2010 the university launched Target, an ambitious research project aimed at revolutionizing the management of very large amounts of data. Under the Target project, prominent research groups and commercial partners will collaborate to develop intelligent information systems capable of transforming huge amounts of raw scientific data into valuable knowledge.

Keeping pace with large-scale research

The amount of data generated by increasingly complex modern research projects is constantly growing. If the University of Groningen is to maintain its reputation as one of Europe's leading research universities, and keep up with the fast pace of research, it must provide scientists with high-performance computational resources that integrate rapid, large-scale data storage and processing.

Wietze Albers, System Administrator at Target, states: "As the amount of data generated by research increases, there is a corresponding need to process that data much more quickly, so that it can be made available for further analysis and computation. This helps scientists to make breakthroughs and complete their research much faster.

"At the same time, we need to guarantee that the Target platform is kept up and running around the clock, not least because Target is intended to support a large international community of researchers and business partners, many of whom work in different time-zones."

Meeting these objectives requires a network infrastructure that combines high bandwidth and redundancy to deliver extremely fast and reliable connectivity for users. The university also aimed to take advantage of the latest networking technology, which would elevate performance and availability while streamlining administration.

Solution Components

Hardware

- IBM® Intelligent Cluster™
- IBM RackSwitch™ G8264
- IBM RackSwitch G8052
- IBM System x® 3550, System x3500, System x3650 M3, System x3690 X5

Software

- IBM General Parallel File System (GPFS™)
- SUSE Linux Enterprise Server

Services

- IBM Global Technology Services®

IBM Business Partner

- Aprycus B.V.

Unique server and storage architecture

The University of Groningen joined forces with IBM Global Technology Services to design and implement a powerful, scalable computing infrastructure supported by high-capacity, high-bandwidth server and storage technology.

The core of the Target platform is an open-standards storage solution called the "test bed", featuring five different storage pools that represent distinct environments. The Target test bed comprises IBM System Storage® tape and disk systems, and currently has a capacity of 10 PB. IBM General Parallel File System (GPFS™) is used to combine the various storage pools into a common global file system.

As it was necessary to stream data to the most suitable storage devices and pools, the IBM team deployed an IBM Intelligent Cluster™ solution comprising 58 IBM System x® servers running SUSE Linux Enterprise Server. This cluster environment includes storage servers, database servers and application servers, and leverages intelligent Intel Xeon processors to deliver high computational density and energy efficiency.

Developing a powerful network infrastructure

The Target infrastructure is distributed over two data centers at the University of Groningen; the locations are interconnected using a fully meshed design that utilizes four 40 GB inter-datacenter connections. These 40 GB links are bundled together using the IBM Virtual Link Aggregation Groups (VLAG) protocol, delivering a fully active backbone network. By negating the need for older protocols such as Spanning Tree, VLAG helps to maximize bandwidth, enabling the backbone network to carry 160 GB of sustained traffic to all attached systems over fully active/active links without the need to block ports.

The IBM System x servers in the high-performance computing cluster leverage high-bandwidth, low-latency Ethernet-based interconnects, all of which have a minimum connection speed of 10 Gb/s. The majority of the servers in the computing cluster feature maximum connection speeds of 20 Gb/s, with 12 servers having connections of up to 40 Gb/s.

Six IBM RackSwitchG8264 switches feature in the cluster environment. The switches are optimized for resource-intensive applications requiring high bandwidth and low latency, making them ideal for use in the Target infrastructure.

"With IBM we have developed a platform that will easily accommodate changing needs in the years to come and one that will offer exciting new opportunities to drive breakthroughs in scientific research."

— Wietze Albers, System Administrator, Target

The University of Groningen has also deployed four 1 GB IBM RackSwitch G8052 switches, which are used in the management layer of the Target infrastructure. The RackSwitch G8052 is an Ethernet switch specifically designed for the data center, providing a virtualized and easy-to-manage network solution.

The IBM System Storage components in the Target test bed are connected to the high-performance computing cluster via a separate combination of Fibre Channel and SAS switches. The university worked closely with IBM Premier Business Partner Aprycus B.V. to implement the networking solution.

Transforming network management with OpenFlow technology

The IBM RackSwitch G8264 switches used in the computing cluster are one of the first 10 GB Ethernet switches to offer the benefits of OpenFlow. This new technology provides an open, standards-based interface to control how data packets are forwarded through the network.

The solution enables network administrators to easily configure and manage virtual networks that control traffic on a "per-flow" basis, and will help to ensure more predictable performance for large transfers of data in the university's complex Target environment.

"We are very interested in taking advantage of OpenFlow technology to build an application-driven network topology," says Albers. "It will offer us an innovative way of managing our network so that, for example, if an application requests more bandwidth, we can make optimizations to our network infrastructure and automatically reconfigure it to help manage the load. This will make it much easier for us to handle the massive amounts of data moving across our network."

Boosting bandwidth levels

The University of Groningen is now able to accommodate the high-performance requirements of scientific applications, thanks to the advanced IBM networking solution developed for the Target infrastructure. With very high bandwidth and low latency, the solution helps to keep vast amounts of data moving at high speed through the Target platform.

"Looking at the bandwidth levels that were possible two years ago, we have increased performance by 50 percent," states Albers. "This is of huge benefit to our users, as they can now complete analyses much faster, whether they involve performing many small computations or processing extremely large data sets. The Target infrastructure can handle it all."

Flexibility for the future

The University of Groningen has gained a modular solution that makes it easy to scale for the future. Next-generation technology such as OpenFlow will support more flexible network management and will make it possible to connect to other open-standards networking solutions in the near future. This will give administrators greater control over traffic management, enabling more efficient data processing.

"With its support for more open standards, the new networking infrastructure from IBM will allow us to adopt a much more flexible approach to network management," notes Albers. "We will be able to deal more easily with the diverse requirements of all the applications supported by Target, which means that we can ensure a higher level of service for our users."

The scalability and flexibility of the Target platform will help to guarantee that the University of Groningen continues to meet the ever-expanding needs of large-scale research.

Albers concludes, "We are constantly redesigning and optimizing our infrastructure to support greater workloads and new users. With IBM we have developed a platform that will easily accommodate changing needs in the years to come and one that will offer exciting new opportunities to drive breakthroughs in scientific research."

For more information

To learn more about IBM System Networking solutions, contact your IBM sales representative or visit: **ibm.com**/systems/networking

To learn more about products, services and solutions from Aprycus, visit: **www.aprycus.com**

© Copyright IBM Corporation 2012

IBM Nederland
Johan Huizingalaan 765
1066 VH Amsterdam

Produced in the Netherlands
December 2012

IBM, the IBM logo, ibm.com, GPFS, Global Technology Services, Intelligent Cluster, RackSwitch, System Storage and System x are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at: ibm.com/legal/copytrade.shtml.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

IBM and Aprycus are separate companies and each is responsible for its own products. Neither IBM nor Aprycus makes any warranties, express or implied, concerning the other's products.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to imply that only IBM's product, program or service may be used. Any functionally equivalent product, program or service may be used instead.

All customer examples cited represent how some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, IBM warranty terms apply.

This publication is for general guidance only.

Photographs may show design models.

Please Recycle

QCC03011-NLEN-00