The worlds of storage and computing combine to form a homogenous environment at the ZIB. Source: Konrad-Zuse-Zentrum, Berlin.

Konrad-Zuse-Zentrum speeds up data storage transfer for high-performance computers with the Cisco Nexus 5020

Konrad Zuse in 1936.
Source: Horst Zuse Archive, Berlin (www.zuse.de).

One of the Z25s, built after 1961, in the ZIB museum.
Source: Konrad-Zuse-Zentrum, Berlin.

No speed limit for Supercomputers

The Konrad-Zuse-Zentrum, also known as the Zuse Institute Berlin (ZIB), has installed two innovative Nexus 5020 switches from Cisco on its high-performance network. For the first time ever, the new components combine storage area network (SAN) and IP-based data traffic in the same unit. By using the Cisco® Nexus 5000, ZIB has very cost effectively increased the speed at which huge volumes of data are transported to and from its in-house supercomputer. This particularly benefits the many scientific institutions that use the supercomputer for tasks such as materials and climate research in cooperation with ZIB. This modernisation project, therefore, promises to improve the potential for innovation for the German research community as a whole.

Powerhouse for cutting-edge research in Northern Germany

Probably everyone in the world benefits from the work of Konrad Zuse, the creator of the binary floating point. His innovation forms the basis for all conventional computers today, from pocket calculators through to high-performance computers. One such supercomputer can be found at the Berlin research institute that bears his name, the Konrad-Zuse-Zentrum für Informationstechnik Berlin, which focuses on algorithmic mathematics and practical information technology. Priority is given to mathematical modelling to describe complex phenomena in science, technology, the economy and the environment. ZIB also develops and tests algorithms for the computer simulation of these models, using theory to help solve practical problems.

ZIB runs the supercomputer together with the Regional Computing Centre for Lower Saxony (RRZN) at Leibniz University, Hanover. Both locations are linked at 10 Gbps - the amount of data that can be transmitted every second at this speed would fill almost a million pages of a printed book. The organisation responsible for this distributed approach is the North German Supercomputing Alliance (HLRN), a scientific cooperation between Berlin, Bremen, Hamburg, Mecklenburg-West

Background

The ZIB is an institution in the Federal State of Berlin which has held an established place in the German research landscape for many years. A key feature of the institute is the close relationship between mathematical and computer science theory and practical applications. As one of the two operators of the HLRN supercomputer system, the ZIB also acts as a service provider for many other scientific institutions.

Challenge

Maximum throughput at the interface between the Cisco SAN and the Ethernet infrastructure is vital for optimum use of the ZIB supercomputer. The institute is also working towards a standardised IP and SAN architecture that will enable it to cost-effectively meet the growing performance requirements in scientific computing.

Solution

Two switches from the new Cisco Nexus 5020 series form a link between the fibre-channel SAN and the IP-based server network. SAN and LAN grow together in this location. Thanks to Fibre- Channel-over-Ethernet, no extra adapters are necessary to connect the server and the SAN. This simplifies the entire infrastructure and makes the server throughput to the data networks much easier to scale by avoiding the bottleneck caused by too few adapter slots in the servers.

Benefits

- Improved performance and reduced complexity
- Greater efficiency because the supercomputer is put to optimum use
- Cost-effective scalability because previous limitations have been overcome
- Improved opportunities for scientific cooperation
- Flexible development of the IP and SAN infrastructure

Nexus 5000, Cisco Systems.

Pomerania, Lower Saxony and Schleswig-Holstein. It would take the entire population of these six federal states around 85 years, with everyone working 12 hours-a-day on complex multiplications, to obtain the computational result that the HLRN supercomputer achieves in one second.

High-speed computers need a high-speed environment

The exceptional computing speed of the 25,000 processor cores can only be used effectively if the surrounding infrastructure can keep up with this tremendous pace. For example, enormous volumes of data from various simulation projects must reach the direct access area of the supercomputer at a fast enough rate, and be swapped out again just as quickly, to make room for the data from the next simulation project. Meanwhile, the online data storage at ZIB has a disk capacity of more than a Petabyte – over a billion Gigabytes. Instead of ordinary hard disks, specially optimised, ultrafast disk systems known as arrays are used. These are linked via a comprehensive SAN with the computers and tape drives of a magnetic tape storage unit. At ZIB, certain parts of the disk storage operate at speeds of 24 Gbps, while the data readout is even faster.

High speed is also a priority in the SAN, which is based on the widely used SAN standard fibre channel (FC). Its most important components are 13 Cisco MDS 9000 Series Intelligent Fabric Switches which together make up 416 x 4gigabit FC ports.

“The interface between the FC SAN and the IP-based server landscape is, to some extent, the bottleneck whose capacity is crucial in determining the performance of the entire system,” says Wolfgang Pyszkalski, manager of ZIB’s IT services department. The throughput at the server-SAN interface is therefore of strategic significance for the further development of the ZIB infrastructure. “The latest milestone in our infrastructure development is the installation of two Cisco Nexus 5020 Series switches. The switches have increased the throughput at this bottleneck between the SAN and the 10-gigabit Ethernet network in a particularly cost-effective way.”

SAN and LAN grow together

The new Cisco access switches overcome the old division between the SAN and the IP network. The Nexus 5020 is the first 10-gigabit switch to make the newly ratified industrial standard Fibre Channel over Ethernet (FCoE) suitable for practical use. FCoE ‘packages’ the fibre channel protocol from the SAN together with the IP protocols from the LAN in such a way that they can be transported via the same Ethernet cables. FCoE has far-reaching consequences for the transformation of today’s data processing centres into flexible service environments with a highly simplified I/O architecture. “The FCoE-capable Cisco Nexus 5020 will considerably reduce the need for us to purchase additional fibre channel adapters. Apart from the direct cost savings, this also means increased scalability of data transfer, as the limited number of adapter slots in a server is a real problem in terms of our requirements,” says Wolfgang Pyszkalski. Because of the Nexus 5020, the SAN and Ethernet servers at the ZIB are also connected via the same unit for the first time.

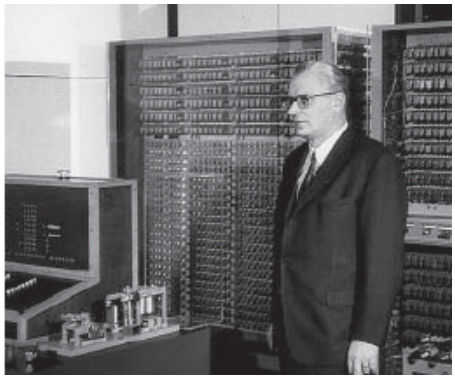

Copy of the Z3 computer with Konrad Zuse (1961). Source: Horst Zuse Archive, Berlin (www.zuse.de).

Konrad Zuse (1910-1995) was a very versatile man: born in Berlin, he achieved fame by building the world's first computer, the legendary Z3, in 1941. After the war, the inventor became an an entrepreneur by founding his own computer company. Zuse relentlessly worked on refining new numerical methods and machines. He wrote books and articles on information technology, society and philosophy. He was also known for his abstract oil paintings and portraits of prominent contemporaries.

Source: Konrad-Zuse Center, Berlin.

The Z22 computer (1958): the computer that introduced electronic data processing (EDP) into German universities and colleges (Statement from Prof. Giloi, Technical University Berlin, 1995). Source: Horst Zuse Archive Berlin (www.zuse.de).

Each server is always connected to both Nexus 5020 switches via separate cables to guarantee maximum failure protection. Each Nexus 5020 is connected to the storage area network via 12 x 4gigabit FC ports. This means a computational bandwidth of 48 Gbps per switch. In contrast, 10-gigabit Ethernet cables leading from the Nexus to two structurally identical Cisco Catalyst 6509 VS-S720 switches form the link to the IP network. The Catalyst switches, also configured to be failsafe and each acting as a virtual switch, serve as both a link to the IP backbone at ZIB, and also as a portal to the outside world – for example, to the German research network X-WiN and the Berlin scientific network BRAIN (Berlin Research Area Information Network).

Unified Fabric within sight

FCoE-capable switches such as the Cisco Nexus 5020 bring the vision of a Unified Fabric within reach for the first time. Unified Fabric in this instance means that, in principle, each server can access both the IP network and the SAN storage media via its Ethernet adapter. "As a result, a large number of adapters, cables and separate control tools will become superfluous," says Wolfgang Pyszkalski. "The standardised Ethernet-SAN fabric considerably reduces the complexity of connections and makes the network more transparent. Consequently, management time and costs are reduced and new components are easier to implement than before." In particular, each of the 40 Ethernet ports of a Nexus 5020 can be connected both by using a conventional 10 Gbps IP adapter and also by using one of the new FCoE adapters. As a result, investment in the new units is protected. They immediately create a 10 Gbps infrastructure at ZIB and offer the facility of a smooth transition to the innovative Unified Fabric by means of FCoE technology.

True to the vision of ZIB's founder, Konrad Zuse, Wolfgang Pyszkalski is already thinking about practical application scenarios. In future, for example, scientific institutions could install additional data security in the ZIB SAN via the Unified Fabric. A prerequisite for inter-organisational SAN usage, however, is that the joint use of individual components can be defined in the SAN, whilst the remaining parts of the networks are securely separated from each other. "This condition has already been met through our SAN, where virtual SANs (VSANs) and Inter-VSAN Routing (IVR) offer precisely this functionality. In fact, these advanced VSAN functions were one of the reasons why we selected the Cisco MDS 9000 technology," says Wolfgang Pyszkalski.

Customer contact details:

Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB)
Takustraße 7
14195 Berlin-Dahlem
Tel: +49 (0)30 84185-0
Fax: +49 (0) 30 84185-125
http://its.zib.de

Cisco Systems GmbH
Kurfürstendamm 21-22
D-10719 Berlin

Cisco Systems GmbH
Neuer Wall 77
D-20354 Hamburg

Cisco Systems GmbH
Hansaallee 249
D-40549 Düsseldorf

Cisco Systems GmbH
Friedrich-Ebert-Allee 67-69
D-53113 Bonn

Cisco Systems GmbH
Ludwig-Erhard-Straße 3
D-65760 Eschborn

Cisco Systems GmbH
Wilhelmsplatz 11
(Herold Center)
D-70182 Stuttgart

Cisco Systems GmbH
Am Söldnermoos 17
D-85399 Hallbergmoos

Tel.: 00800-9999-0522
www.cisco.de

For technical assistance regarding the Cisco product range or questions on your network design, please contact the Cisco Technical Helpdesk on **0800-9999-0522** or e-mail information@external.cisco.com.