National Centre for Nuclear Research

Highlights

Challenge

Polish National Centre for Nuclear Research and its Świerk Computing Centre needed to improve its core network infrastructure to facilitate the upgrade of its supercomputing resources and provide a more efficient foundation for further growth.

Solutions

- Arista 7050 Series Switches
- Arista EOS®
- Coraid EtherDrive

Results

- Wire-speed performance on all switching ports
- Significant reduction in latency
- Dramatically lower energy consumption
- Switch to MLAG doubles bandwidth capacity
- HPC application running directly within the switch boost computational performance
- Advanced admin tools (XMPP, JSON RPC) reduce management complexity and speed up configuration changes

Arista helps National Centre for Nuclear Research build a low latency and high performance network infrastructure to support supercomputing excellence

Narodowe Centrum Badań Jądrowych
National Centre for Nuclear Research
Świerk

The Polish National Centre for Nuclear Research provides a catalyst for pure research and numerous practical applications across science and the wider economy. In response to demands from the Polish national energy sector, the Centre established the Świerk Computing Centre and began a major program to upgrade its critical infrastructure to increase performance with the ability to scale further over the next decade. Through switching to Arista, Świerk has dramatically increased its uplink connections bandwidth, while reducing latency and built the foundation to scale its IT resources in line with new groundbreaking scientific and commercial projects.

Project Background

With over 1000 employees, the National Centre for Nuclear Research is one of the oldest and the largest research institutes in Poland. Since its inception in 1955, it has established a world-class combination of pure research and numerous practical applications for science and the wider economy. NCNR is one of only a handful of world suppliers of particle accelerators for industry and medicine as well as an operator of the only nuclear reactor in Poland and a key Technical Support Organisation (TSO) for the national nuclear power program. The institute cooperates with leading universities and R&D centres including close links with CERN, the largest scientific laboratory ever built and at the same time offers its own unique research infrastructure for domestic and foreign scientists.

A crucial part of NCNR's resources is the Świerk Computing Centre - one of the largest supercomputers in Poland and the region. The centre was established in 2009 with aid from the European Regional Development Fund (Innovative Economy Operational Program) and through co-financing by the Polish Ministry of Science and Higher Education. With a total cost of approximately PLN 98 million, the main goal of the supercomputing resource, together with a group of experts and scientists, is to provide the information technology support needed for the development of the Polish nuclear and conventional energy sector. The compute cluster is being used to monitor and simulate radiation hazards, emergency management, as well as calculate for the design, installation and optimization of power equipment and power distribution. Furthermore, Świerk Computing Centre runs a number of its own research and development projects related to the safety and operation of nuclear reactors, chemical analysis, the physics of fundamental interactions and astrophysics, as well as work for the design and construction of medical equipment.

For such exacting tasks, the centre is continually looking at technologies to help it harness its vast computing capability and mitigate any technical hurdles to help it deliver its world class research capabilities. Although currently undergoing a validation process, the installation is likely to be categorised as one of the top 500 most powerful supercomputers in the world and certainly in the top 3 in Poland. The strategic goal is to build a unique science and technology competence centre to support the development of the Polish energy sector and all related fields of science. To meet this aim, an architecture refresh would be required to allow the cluster to scale as needed, but with particular care to avoid any bottlenecks that could impact future performance. One such area of particular concern was the core networking technology that would underpin the centre for the next decade of expansion.

Solution

NCNR specified an open tender to a number of suppliers for both compute and networking upgrades with the technical team at the centre scrutinising every response. As Piotr Szwajkowski, Head of Network Security Department at NCNR explains, "It should be noted that the bar regarding the technological requirements was placed very high. Building HPC clusters requires a sufficiently large bandwidth of individual network connections as well as very small delays in data transmission. Our HPC cluster will be expanded in the future, so when choosing an appropriate network solution, a very important aspect was also its scalability."

As traditional Spanning Tree Protocol (STP) solutions were considered unsuitable, all vendors were obliged to offer Multi Chassis Link Aggregation (MLAG) functionality. "Due to the MLAG approach and Spine-and-Leaf architecture implemented in Arista network devices, we were able to opt out of the traditional L2 loop-free mechanism offered by STP and use the capabilities of installed devices in more efficient ways," Szwajkowski explains, "The MLAG protocol and Spine-and-Leaf architecture also offers a higher availability level as there is no single point of failure in the whole network."

This viewpoint was shared by Adam Padée, Head of the Computing Infrastructure Team at the Centre, "When we started building our installation, the number of solutions on the market that offered MLAG or similar functionality was very limited. The advantage of the solution offered by Arista is a combination of MLAG functionality with very low packet latency, which is a very important feature when it comes to distributed calculations."

"In addition, these switches are characterized by extremely low energy consumption which was up to three times less than competitors offering similar functionality, so it is an attractive feature when taking into account the total energy consumption of the supercomputer," Padée adds. Another feature that was of particular interest was the ability to create applications scripts directly within the Arista switches to perform complex cluster management processes with much greater efficiency and fewer performance overheads.

Conclusion

With Arista offering multiple technical, power, cost per port and space benefits over all of the rival bids, the Centre decided to implement the new network architecture. The core was based around four Arista 7050Q series 40G switches offering wire speed layer 2/3/4 performance each with high density 40GbE in a compact 1RU chassis. The top of each of the four compute racks included Arista 7050S series switches with 64 ports of 10GbE in a compact 1RU chassis. All switches use the same Arista EOS, a modular switch operating system with a unique state sharing architecture that cleanly separates switch state from protocol processing and application logic. Built on top of a standard Linux kernel, all EOS processes run in their own protected memory space and exchange state through an in-memory database. This multi-process state sharing architecture provides the foundation for smart system upgrades (SSU) and self-healing resiliency.

Following a two week installation process, involving teams from Arista, local IT partner Format and the Świerk Computing Centre IT department, Szwajkowski stated, "The implemented Arista network solution in Świerk Computing Centre has uplink connections bandwidth of 320Gbps and extremely low port-to-port latency of less than 1us. This allows us to fully utilize installed computing devices within the HPC cluster."

The move to Arista has delivered lower energy consumption then other options and OPEX has been further reduced through a reduction in management overheads. The ability to run certain applications within the switches combined with powerful scripting capabilities has significantly reduced the amount of time spent needed for cluster administration.

The complete upgrade to the main server room at the centre now includes 1024 new generation 10 core processors supported by 73TB of RAM. In this first phase, the theoretical efficiency of the compute resources has increased from 17 to 281 teraflops (trillion floating point operations per second). By May 2015, the main computer cluster efficiency in the Świerk Computer Centre will reach 500 teraflops supported by the new Arista network infrastructure.

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062