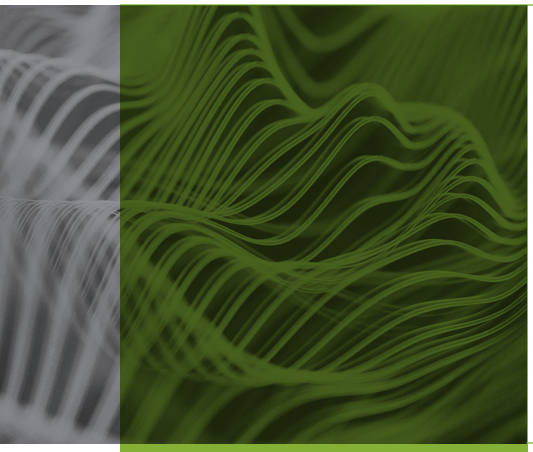

supplier

CERN PEERS INTO THE ORIGINS OF THE UNIVERSE WITH A NETWORK BUILT TO HANDLE A DATA DELUGE

Summary

Company:
CERN

Industry:
Scientific Research

Business Challenges:
Prepare for the third run of Large Hadron Collider experiments, which will produce unprecedented volumes of data.

Technology Solution:
- *QFX Series Switches*
- *EX9200 line of Ethernet Switches*
- *The Linux Foundation Tungsten Fabric*

Business Results:
- *Delivered a resilient network to support analysis of petabytes of particle physics data*
- *Enabled a programmable network for increased flexibility and simpler operations*
- *Automated configuration and management of 4400 routers and switches*

What's the universe made of? How does it work? Answering the big questions about the universe is the mission of European Organization for Nuclear Research (CERN). Home to the world-renowned Large Hadron Collider (LHC), CERN is advancing the boundaries of human knowledge through breakthrough research in fundamental physics. CERN chose Juniper Networks for powerful, high-density switching that is helping unlock the mysteries of the universe.

The Higgs boson is a cornerstone of particle physics. Discovered at CERN, researchers around the world continue to study the way the Higgs boson decays or transforms. Through the experiments run on the LHC, scientists have discovered exotic new particles and observed never-before-seen phenomena.

After a second outstanding three-year run of achievements, CERN shut down the LHC at the end of 2018, with the goal of restarting in 2021 with an even-more powerful collider to further expand our understanding of the universe. In preparation for Run 3, the LHC will undergo a significant increase in instantaneous luminosity to 1.5 times its current value.

"When the LHC experiments are turned back on, the amount of data the detectors will send to the data center will be significantly higher," says Vincent Ducret, network engineer at CERN. "Two of the experiments will go up to terabits of data per second, for example."

CERN's IT team has 1.5 years to build a network that supports the expansion of our knowledge of fundamental physics. CERN chose Juniper Networks for the technical network that supports the LHC operations and experiments as well as its data center and backbone. All the switches and routers run the same Juniper Networks® Junos® operating system, which simplifies the configuration and management of hundreds of routers and switches across CERN's campus. CERN also leverages Tungsten Fabric, formerly OpenContrail® and now part of The Linux Foundation, as part of its OpenStack cloud.

"Junos operating system is programmable and has a very good API. Because all of our Juniper devices run Junos OS, it facilitates our task of supporting different models, including the QFX10000 and EX9200 switches.

- Stefan Stancu, network and software engineer, CERN

At the Heart of Scientific Discovery

CERN needed to increase the capacity of its data centers, including servers, storage, and networking, to fuel the next wave of discovery. The network needed to be faster and more redundant. The network also had to be programmable, enabling automated configuration and management, and adapt easily to changing requirements. "The goal wasn't only more capacity," says Ducret. "We needed the network path to be more dynamic. We wanted the ability to use technologies like VXLAN for greater flexibility in connecting servers over the network."

Detectors—akin to highly specialized digital cameras—are trained on the proton collisions in the LHC chambers; the results are monitored in the hope of discovering new particles and new reactions between particles. Collisions generate about one petabyte of data per second. This data is filtered by the experiments, which keep only what is considered "interesting." The data is then aggregated, reconstructed, and archived to long-term storage for later study. Copies of the data are sent to other large-scale data centers, forming the Worldwide LHC Computing Grid (WLCG).

"The goal wasn't only more capacity. We needed the network path to be more dynamic. We wanted the ability to use technologies like VXLAN for greater flexibility in connecting servers over the network."

- Vincent Ducret, network engineer, CERN

The Geneva data center is massive. It supports all of CERN's scientific projects and experiments and administrative systems—everything from spotting camera-shy pentaquarks and charm mesons to everyday videoconferencing and payroll. The data center is home to more than 15,000 servers and 260,000 processor cores and has a storage capacity of 280 petabytes. More than 340 petabytes of data are archived in tape libraries.

The Juniper Networks QFX10008 Switch is used for the data center core network. Designed for the most demanding data center and multicloud environments, the QFX10000 line of Switches uses Juniper's custom silicon Q5 ASICs for high-performance, versatility, scale, and capacity for long-term investment protection. The QFX10008 delivers on CERN's requirements for fast, reliable connectivity and 100 Gbps port density to support the inevitable traffic growth. The core network can be scaled from 40 Gbps to 100 Gbps and above with the Juniper Networks QFX Series Switches.

The QFX5120 and QFX5200 Switches, high-density 25GbE and 100GbE fixed-configuration platforms, provide full redundancy in top-of-rack deployments, an enhancement in availability. High-density switching is essential for flexibility and cost-effectiveness and reduces the rack space and cooling requirements for future expansions.

Serving at the data center edge, the QFX10002 Switch connects CERN's data centers to the networks that distribute data to research institutions around the world for compute and analysis. QFX10002 fixed-configuration switches are very dense, highly scalable platforms in a compact 2 U form factor. This grid provides about 80 percent of the compute power. In 12 months of the last LHC run, 370 petabytes of data have been moved through these networks.

Open source plays a vital role in meeting the research demands of collaborators around the world. CERN runs a private OpenStack cloud with more than 280,000 cores and 8000 hypervisors. It has deployed a new region which leverages SDN for improved workload mobility across the data center, floating IPs, and IP-based server load balancing. Tungsten Fabric, a community-led, open-source project and the basis for Juniper Networks Contrail® Platform, provides the network control plane in this region. CERN is currently investigating the use of Virtual Extensible LAN (VXLAN).

Capturing the Ephemera of the Big Bang

ATLAS, one of the four main experiments being conducted, is designed to observe up to 1.7 billion proton-proton collisions per second and produces a data volume of more than 7.5 terabytes per second. Only some of the observed events are sufficiently interesting to lead to new discoveries, so the data flow is reduced to a manageable level. The ATLAS Trigger and Data Acquisition System (TDAQ) handles this filtering, as well as data collection and infrastructure monitoring.

The QFX5200 Ethernet Switch provides high-throughput connectivity for the servers, while a pair of Juniper Networks QFX10016 Switches configured with multichassis link aggregation groups (MC-LAGs) provide redundancy, load balancing, and a loop-free Layer 2 network without the overhead of the Spanning Tree Protocol. Multiple virtual router instances with different class-of-service policies allow different traffic types, such as management, control, data, monitoring, and simulation, to be treated differently. The QFX5200 and QFX10016 chassis serve the detector read-out channels on ATLAS.

Automation Makes Light Work

CERN's Technical Network, which supports the operations of the accelerator as well as the monitoring and safety systems, is the most critical element of its data infrastructure. Connecting more than 11,000 devices, the architecture emphasizes robustness, reliability, and scalability. CERN deployed the Juniper Networks

EX9200 line of Ethernet Switches, a family of programmable, flexible, and scalable switches, for its Technical Network.

Infrastructure security is a major concern, and scientists now have a broader choice of security policies they can use to protect their systems. Access between the Technical Network and the General Purpose Network is governed by extensive access control lists (ACLs). The fact that Juniper's switching is able to handle ACLs of this size is allowing the network engineering team to make current operations more efficient while designing a future solution for firewall filtering.

Leveraging open source tools, the network engineering team has developed software to produce the specific device configurations based on a database that models the desired network state. Juniper devices are configured based on these generated configurations.

Automation is critical to configure, deploy, and manage a very large-scale, multivendor network. "All our network is provisioned automatically," says Stefan Stancu, network and software engineer in the IT department at CERN. "A service technician can change a switch, and our software tools then fully provision its configuration in a single command. With hundreds of routers and thousands of switches, manual configuration would be prone to human errors. Furthermore, we automatically update the contents of all routers' access control lists on a regular basis, and we rely on Juniper's advanced configuration interface to detect changes and apply the updates only when necessary."

"Junos operating system is programmable and has a very good API," Stancu continues. "Because all of our Juniper devices run Junos OS, it facilitates our task of supporting different models, including the QFX10000 and EX9200 switches. Support for open standards through programmatic interfaces is also important. We use our own, internally developed tool to configure and manage the network equipment. This allows us to have a vendor-agnostic user interface and our operational workflow remains the same regardless of which vendor is used. However, under the hood, we need to program the interaction with each different vendor, and in this respect, Juniper facilitates our task by having a clear programmable configuration API, which is virtually identical for all Juniper platforms."

Preparing for Even Higher Luminosity

CERN's experiments will continue to evolve, and the network engineering team sees the challenges ahead. The High-Luminosity LHC, to be commissioned in 2024, will increase the precision of the ATLAS and CMS detectors. The requirements for data and computing will grow even more dramatically, with rates expected to be around 500 petabytes per year. CERN openlab and other R&D programs are investigating how to evolve computing models to address these needs.

For More Information

To find out more about Juniper Networks products and solutions, please visit www.juniper.net.

About Juniper Networks

Juniper Networks brings simplicity to networking with products, solutions and services that connect the world. Through engineering innovation, we remove the constraints and complexities of networking in the cloud era to solve the toughest challenges our customers and partners face daily. At Juniper Networks, we believe that the network is a resource for sharing knowledge and human advancement that changes the world. We are committed to imagining groundbreaking ways to deliver automated, scalable and secure networks to move at the speed of business.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: +31.0.207.125.700
Fax: +31.0.207.125.701