

---

**Author's Name:** Paula Holland

**Content Owner/Sponsor's Name:** Sandra Livinghouse

**Marcom:** Sandra Livinghouse

**Editor:** Jonas Weisel

## Executive Summary

**Customer Name:** Sandia National Laboratories

**Industry:** Government/Defense

**Location:** Albuquerque, NM

**Number of Employees:** 8500

**Headline:** Sandia National Laboratories Innovates Research & Development and Testing with High-Performance Computing Solution

**Subhead:** Sandia selects Cisco's InfiniBand Server Switches to create the largest high-performance cluster in the world to meet defense simulations and research computing demands

### Business Challenge:

- Needed to create a high-performance cluster for research and development of national defense, energy and environment projects.
- Desired class of high performance cluster had never been built before.
- Needed a solution that would meet technical requirements that included scalability, high performance, reliability and ease of troubleshooting.

### Network Solution:

Create a high performance cluster employing Cisco InfiniBand Server Switches and software for acceleration and high performance computing.

### Business Value:

- Higher performance than general proprietary system for a significantly lower cost
- Perform simulations and parallel jobs without compromising overall cluster performance
- HPC solution meets current needs and serves as a template for future high performance computing and reliability.

### Pull Quote:

“Sandia’s goal was to transform how we do engineering and provide greater processing capacity. With its 4,480 commodity compute servers linked with an Infiniband message passing interconnect, Thunderbird is the largest cluster of its type in the world. Infiniband is widely regarded as one of the most attractive commodity interconnect technologies, because of its high bandwidth, low latency and low cost.

Thunderbird’s achievement was a joint venture involving Sandia, Dell, and Cisco Systems, Inc. As an active developer in the OFED and OpenMPI projects, Cisco’s support and contributions were invaluable. Cisco engineers were on site at Sandia to assist with monitoring, diagnosing, and fine-tuning Thunderbird’s performance. For the High Performance Computing community, this environment allows for more memory

per node to be available for parallel jobs at runtime. The results have also shown an increase in reliability and scalability of users' jobs.

– Sophia Corwell, Capacity, Computing and Visualization Project Leader, Sandia Labs

**First and Last Name:** Sophia Corwell

**Job Title:** Capacity Computing and Visualization Project Leader, Scientific Computing

**Company Name:** Sandia National Laboratories

## BUSINESS CHALLENGE

Sandia National Laboratories was founded in 1949 in response to the changing national security needs of post World War II United States. Today, Sandia is a multi-program lab that primarily conducts research and development for national defense, energy, and environment projects, and is dedicated to solving national and global threats to peace and freedom for the 300+ million that reside in the United States. In addition to R&D, Sandia is also responsible for nuclear stockpile maintenance and weapons simulation, and relies on a variety of high end computing technologies within the organization's IT infrastructure.

In their quest to stay on the cutting edge of computing, scientists at Sandia made a long term investment to create a cluster to support the greatest computing processing capacity possible in a manner that had never been done before. They decided to build a high performance Linux cluster for the purpose of transforming engineering to conduct a broad range of weapons simulations, including nuclear stockpile maintenance, atomistic scale-to-device modeling of radiation effects on semiconductor electronics, assessing weapon-response safety in extreme thermal and impact environments, and weapons simulation to quantifying uncertainties in performance. According to Sophia Corwell, Capacity Computing and Visualization Project Leader for Sandia Labs, creating the new, high performance cluster presented a challenge for the lab, as they had never embarked upon a project of this scale.

“We were somewhat pioneers in building this type of high performance cluster,” said Corwell. “As far as we knew, this was the first time anyone had ever built a cluster of this magnitude in any field in terms of existing technologies. We knew we would need high performance class switching capabilities for the massive 8,960 processors that would comprise the Thunderbird cluster.”

As scientists began closely examining high end network switching solutions, they understood that a fabric of this class required a solution that would give them unprecedeted bandwidth, superior latency and high availability to deliver greater agility to the cluster. Specific technical requirements included the following:

- Scalability – given the complex configuration of the cluster, the team needed a high performance switching solution that would scale dramatically without degrading application performance and without placing limitations on the 4480 nodes.
- Performance – with a cluster the size of Thunderbird, the team needed the InfiniBand switching capabilities to dramatically reduce latency while increasing bandwidth and speed without overheating.
- Reliability – needed a highly reliable fabric along with hardware and software solution that would have a high mean time between failures, and could continue to perform well in the presence of individual component failures.

- Ease of Troubleshooting – with a cluster as large as Thunderbird, it was important for engineers to be able to quickly and accurately identify a problem.

## NETWORK SOLUTION

Sandia required a high-end switching solution that was stable and would maximize system availability. In addition, they needed something that would scale Message Passage Interface (MPI) and applications to heights they had not been able to achieve previously. Sandia decided on an architecture comprised of 276 Cisco SFS 7000 Series InfiniBand Server Switches as leaf switches that connected with Dell PowerEdge 1850 1U Servers, and eight Cisco InfiniBand 7024 Server Switches that served as core switches.

Additionally, the network used two dedicated 1U servers to running the Cisco High performance Subnet Manager (HSM) in a redundant configuration.

The host software stack consisted of Linux 2.6.xx with the Open Fabrics Enterprise Distribution InfiniBand drivers and Open MPI as the MPI library. Through aggressive and cutting-edge research techniques, Open MPI dramatically reduced the amount of memory required for parallel applications as compared to prior solutions. Open MPI also allowed Sandia to scale their applications to larger problem sizes at finer resolution than they had ever run before on the same hardware. According to Corwell, this precision increase also improved overall computing efficiency, a fact revealed by the 14.7 teraflop increase in Thunderbird's Top500 Linpack rating from 2005 to 2006.

According to Corwell, the high performance, scalability, robustness and flexibility of the Cisco InfiniBand solution enabled Sandia to perform a variety of computations and perform simulations, and run parallel jobs without compromising the cluster performance.

“We perform a wide variety of computational tasks including a broad range of weapons simulations to assess weapon-response safety in extreme thermal and impact environments,” says Corwell. “The computational demands for these types of undertakings goes way beyond the typical network switch and after careful examination of available high-performance solutions, our search led us to Cisco InfiniBand. The fast switching capabilities of InfiniBand allowed for more memory per node to be available for parallel jobs at runtime, as well as an increase in reliability and scalability of users’ jobs. The 4,480 Dell servers linked with the Open MPI is the largest cluster type in the world.”

According to Corwell, Cisco InfiniBand helped enable Sandia’s software-stack environment, allowing for more memory per node to be available for parallel jobs at runtime. InfiniBand also helped improve performance, reliability and scalability of users’ jobs.

During bringup of the Thunderbird cluster, Cisco implemented a modular bringup process, which allowed multiple teams to find and correct installation issues across many parts of the cluster in parallel, and Cisco and Sandia jointly enhanced Cisco’s existing network management tools to aid in quickly locating problem components. The Cisco SM’s reporting of networking events has been extremely useful during both bringup and ongoing maintenance of the network. The SM’s ability to generate log events identifying individual components when error counters exceed thresholds, and when ports in the network change state is in constant use, and Sandia has tied this information into their job scheduler to keep the pool of available nodes consistently up to date. The HSM’s db-synchronization feature has also been instrumental in

improving network reliability. There have been very few SM failures in the Thunderbird network, but the ability to service one SM node while the other takes over has been used numerous times.

“The Thunderbird cluster achieved an overall 38 percent performance improvement from the last year’s Linpack results due to tweaking and fine tuning the software stack,” says Corwell. “Changes we made that contributed to the results were due to switching to the OpenFabrics Enterprise Distribution (OFED) InfiniBand driver stack and Open MPI in which the Cisco team played a key role. We also achieved an 18.5 percent increase in per node efficiency and applications were reporting roughly 18-20 percent increase in speed/performance due to the switching capabilities.”

Corwell sites the limitations of the nodes that had previously presented a challenge were also addressed once the switching capabilities had been optimized in the software stack.

“Initially, we were not able to utilize two processors per node and could only use one processor per job,” says Corwell. “After implementing changes to the software stack, we were able to allow users to run jobs at 1024 nodes at once as opposed to just 512 nodes.

## **BUSINESS RESULTS**

Corwell says the overall stability, systems uptime, scalability of the MPI and the overall ability to scale applications to heights they could not achieve before are key benefits from implementing the Cisco InfiniBand, OpenFabric Enterprise Distribution (OFED) and OpenMPI.

In addition to disaster scenarios and defense planning, Corwell credits the cluster to assisting Sandia’s development in other areas as well.

“InfiniBand has dramatically increased our range of scope, allowing our application designers to perform in a way that had previously been impossible,” says Corwell. “The Thunderbird cluster is now able to scale to perform tasks such as homeland security, modeling biological cells, difference codes/projects scientists are involved in to drive new materials. We also partner with energy system for oil companies and creative solutions for other fields such as modeling climate changes. The cluster powered by InfiniBand helps us drive new ways to leverage computing.”

Corwell attributes the improvement in the cluster’s performance to the propelling of its switch to OpenFabric Enterprise Distribution (OFED), a Linux-based open-source software stack qualified by the OpenFabrics Alliance to operate with multi-vendor InfiniBand hardware and Open MPI, an open-source consortium that has created a portable, high performance implementation of the Message Passing Interface (MPI) standard.

“Our achievement with the cluster was a result of a joint venture involving Sandia and Cisco,” says Corwell. “Cisco was instrumental in helping us maintain the infrastructure and helped with our software builds and network management of such a large system on a day to day basis. They were especially helpful with regards to helping with ensuring application scalability through the Open MPI software that Cisco co-developed.”

## **NEXT STEPS**

As Sandia continues to develop their computing fabric, Corwell believes the achievements they have realized with the help of Cisco will benefit the entire high-performance-computing community. As a result of this project and its impressive results, Sandia moves forward in capacity computing leadership across the science and technology community here in the US.

“Without InfiniBand, we would not have had the necessary high performance switching capabilities for the cluster,” says Corwell. “The collaboration with Cisco for Open MPI and OFED projects has greatly impacted Thunderbird. We are so impressed with the cluster’s overall perform, we’ll continue strengthening our relationship with the vendors involved in the Thunderbird project.”

## **For More Information**

Find out more about Cisco InfiniBand Server Switches and solutions, please visit <http://www.cisco.com/go/hpc>.

## **Product List**

Cisco Application and Storage Networking Services:

- Cisco SFS 7000 Series InfiniBand Server Switches