SCinet Architecture Featured at SC18: The International Conference for High Performance Computing, Networking, Storage and Analysis

Lance Hutchinson
Sandia National Laboratories
Albuquerque, NM, USA
lvhutch@sandia.gov

Julian Hammer
Friedrich-Alexander-Universität
Erlangen-Nürnberg
julian.hammer@fau.de

Corby Schmitz
Argonne National Laboratory
Lemont, IL, USA
cschmitz@anl.gov

***Abstract*—Each year SCinet builds a temporary network at the forefront of technological development in collaboration with research institutions, network equipment vendors and wide area network (WAN) providers, to support technical demonstrations of participating exhibitors and SCinet supporting organizations during the Supercomputing (SC) conference.**

For SC18, 145 exhibitor connections are provided by SCinet with up to 100 Gigabit per second each and a combined bandwidth of 5.13 Terabits per second. SCinet itself is connected to wide area networks with a record breaking 4.02 Terabits per second through six providers.

Planning for this endeavor began a year before and implementation took one month, the network will be in operation for one week and tear down takes place in one single day. 2018 is the 26th year of SCinet, with 225 volunteers out of 85 organizations and 52 million USD of support from vendors, institutions and WAN providers.

Here we present the 2018 architecture, as will be operated Nov 11-16, 2018 during the SC conference at the Kay Bailey Hutchison Convention Center Dallas, Texas. This year SCinet included 400 Gigabit ethernet transports for the first time into its architecture.

I. Introduction

SCinet is a purpose-built network for the International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing or SC). Each year, SCinet brings to life a high-capacity network that supports applications and experiments that are the hallmark of the SC conference. The network links the convention center to research and commercial networks around the world. This experimental network serves as a platform for exhibitors to demonstrate bleeding edge technologies, research institutions to showcase their latest infrastructure, test experimental applications and apply the latest developments in network research, as well as evaluating multi-vendor interoperability.

The Network Research Exhibition (NRE) represents coordinated demos taking place at SC, both on the show-floor and elsewhere with live results visualized on the show-floor. The NRE demos utilize bandwidth across nearly all of the links brought to the show-floor, with special configurations required to facilitate virtual point-to-point connections between remote research institutions/facilities and show-floor booths. Of the 4 Tbps of bandwidth brought to SC18, better than 90% is directly tied to the experiments and demonstrations articulated in the 36 accepted NRE submissions. This includes, but is not limited to:

- Multiple 200 Gbps and 100 Gbps Data Transfer Nodes showcasing enhanced transfer technology for moving various datasets across the world.
- A 400 Gigabit Ethernet switching infrastructure highlighting cutting-edge networking equipment operating at next-generation speeds and the interoperation of this equipment with legacy networking equipment.
- Multiple security-related demonstrations attempting to analyze and protect high-speed networks at 100 Gbps through data organization and orchestration.

Volunteers from academia, research, government and industry work together to design and deliver the SCinet infrastructure. Industry vendors and carriers donate millions of dollars in equipment, services and people needed to build and support the local and wide area networks. Planning begins more than a year in advance of each SC conference and culminates in a high-intensity installation staging and setup phase in the days leading up to the conference.

II. Architecture

The attached SCinet architecture for SC18 illustrates the overall layout of the network. We will give a short description on the illustration, followed by a brief discussion of this years highlights.

At the top, the clouds represent the organizations that handle remote termination points for the wide area network (WAN) circuits which provide all uplink bandwidth to SCinet — some of them thousand of miles away (e.g., C37 to Singapore). On the bottom are the booth end points located on the SC show-floor with the according booth numbers, connected via fiber-optic cables. All booth end points are connected through distributed network operation centers (DNOCs) to the central network operation center (NOC). The NOC and the DNOCS contain exclusively equipment operated by SCinet: green

This technical report was compiled to allow citation of the SC18 SCinet architecture. Please cite using DOI 10.5281/zenodo.1470503.

boxes represent routers, blue are layer one and two devices (e.g., switches, media converters or data center interconnects), all other boxes are servers, appliances and test equipment. The right hand side makes up the conference network, providing wired and wireless connectivity to conference participants, speakers and committee, as well as basic network services for network deployment and management (e.g., DHCP, NTP). Links between devices are represented by lines, with line speeds and technology indicated with color.

The SC18 architecture showcases many technologies throughout various parts of the network. Data Center Interconnect (DCI) technology is utilized to consolidate the number of fiber-optic cables needed to connect the DNOCs back to the SCinet core network. These DNOCs, spread across the exhibit floor, aggregate booth end point connections into multiple 100 Gigabit uplinks. Large-scale participants through the SCinet Network Research Exhibition have been provided a Booth NOC (BNOC) to demonstrate delivering circuit termination closer to the bandwidth consumer, eliminating the need for dedicated fiber for each circuit. Each DNOC and core NOC was equipped with a Software-defined Network (SDN) router to seamlessly provision multiple hardware platforms through the Faucet vendor-independent SDN controller. Inclusion of a layer-0 optical switch into the SCinet core provided flexibility for the interconnect between inbound WAN circuits, test equipment, network taps, and core routers.

III. Conclusion & Future Work

This year's SCinet architecture was designed to meet the needs of commercial and academic network research communities, exhibitors, and conference attendees during SC18. During the upcoming week of operation, SCinet will evaluate the implementation of this architecture through engaging with the user communities and collecting performance measurements and statistics. We will pay specific attention to the advantages and disadvantages of the SDN provisioning and optical switching technology implementations to more effectively utilize existing limited resources. This will allow us to evolve and design next year's iteration of the SCinet architecture based on past experience, best practices and upcoming research.

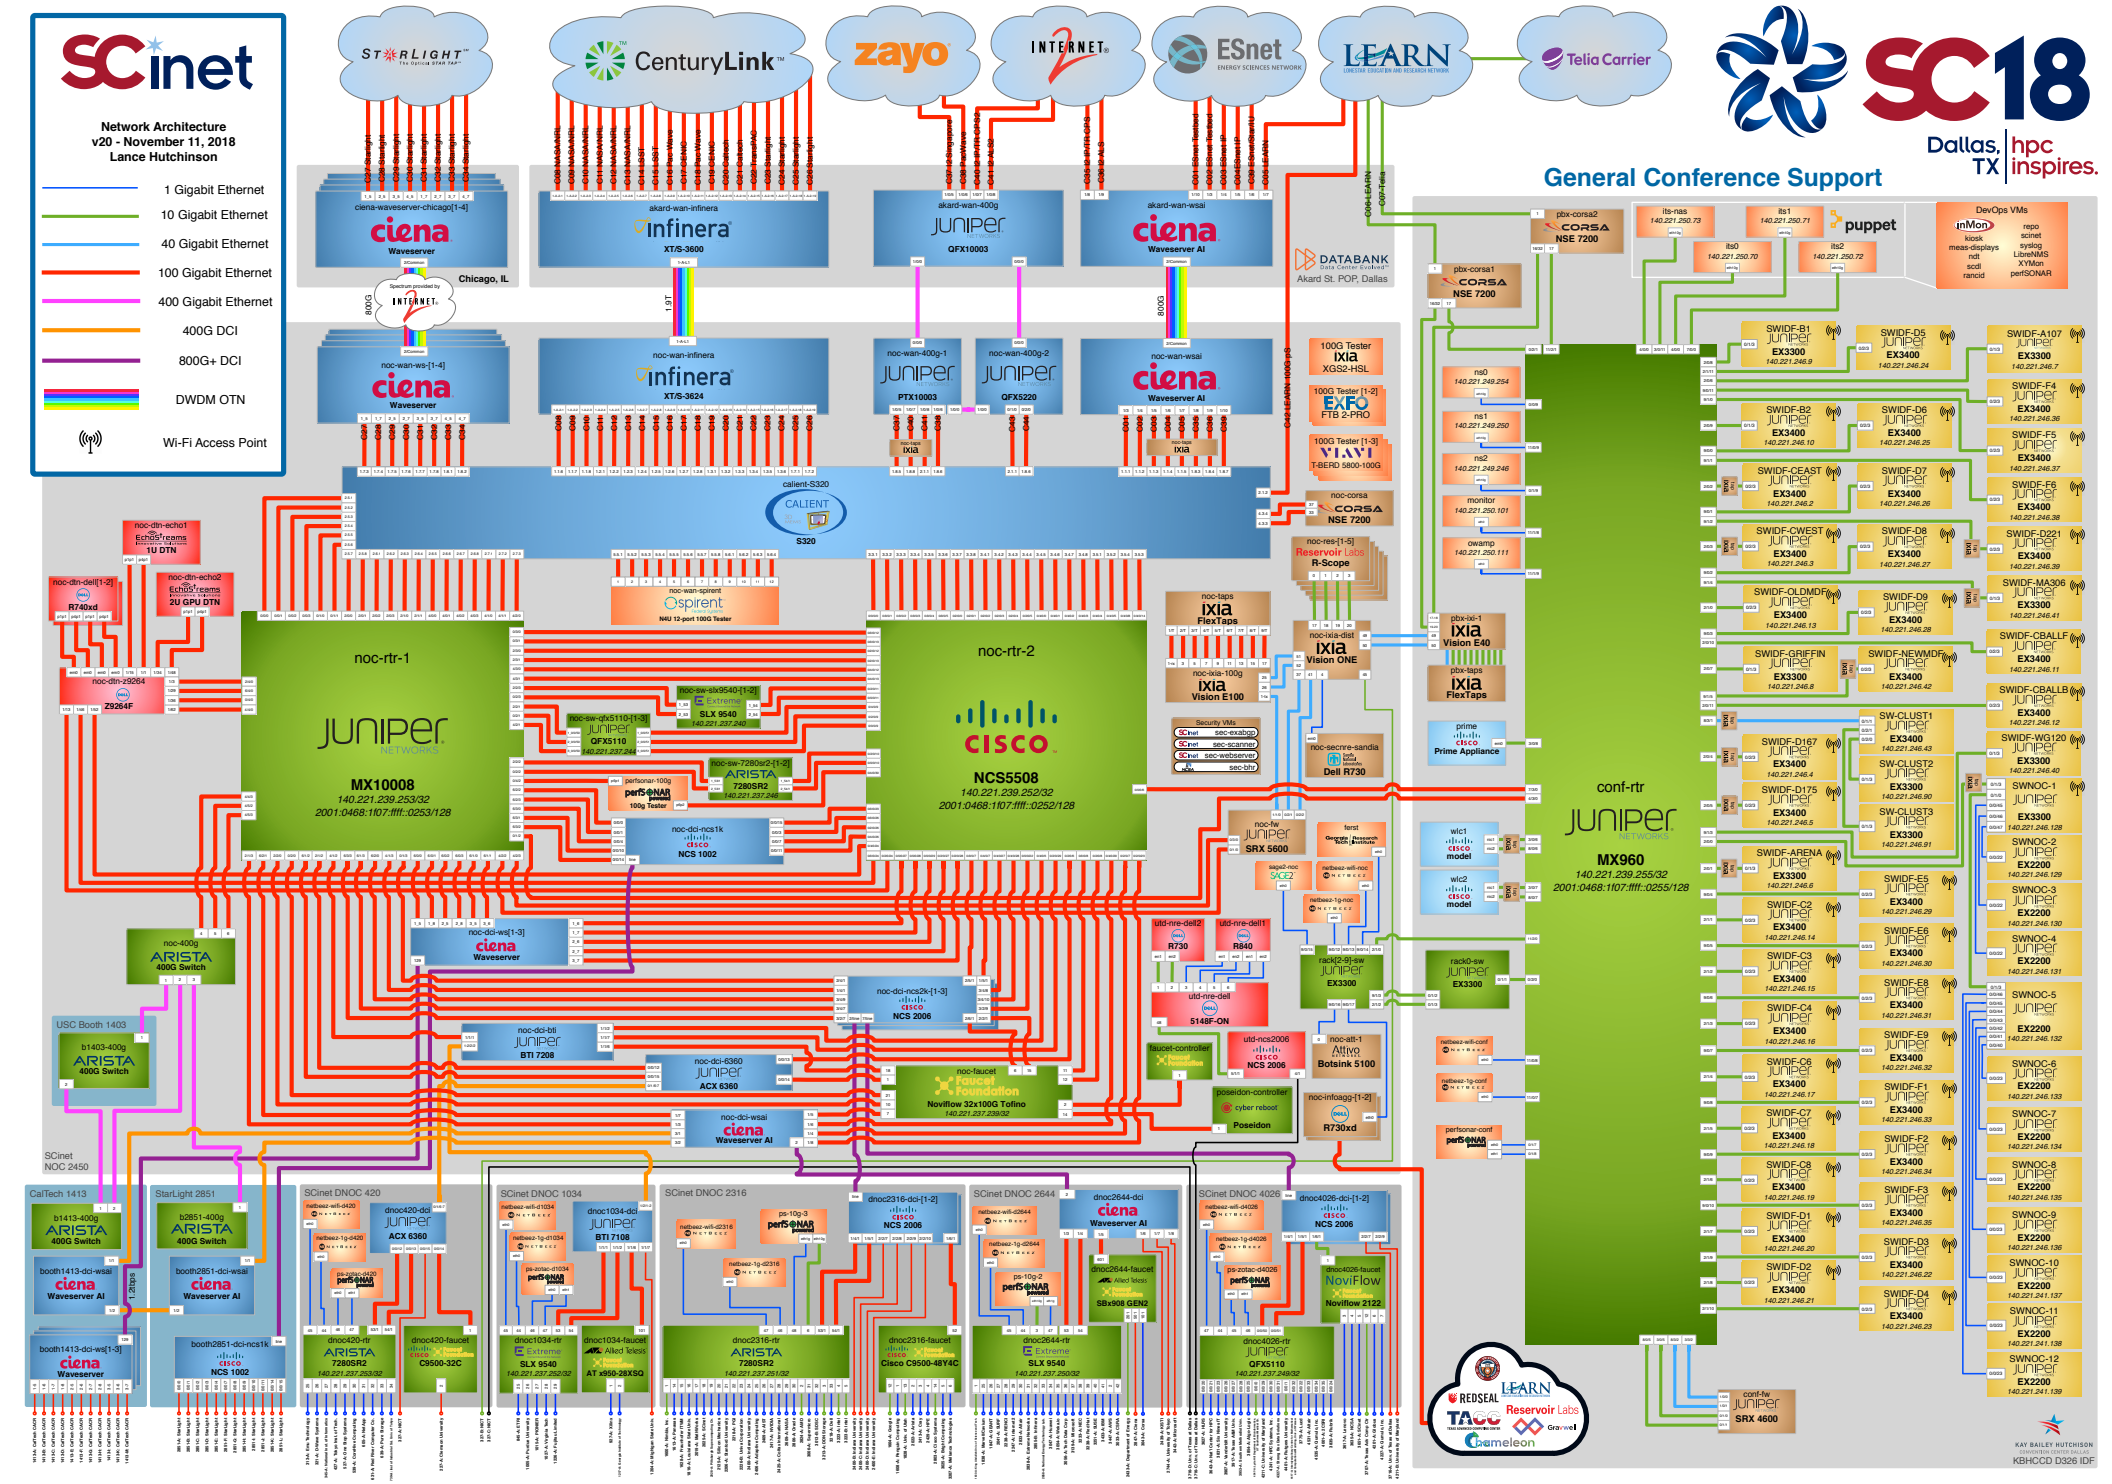
SCinet
Network Architecture
v20 - November 11, 2018
Lance Hutchinson
1 Gigabit Ethernet
10 Gigabit Ethernet
40 Gigabit Ethernet
100 Gigabit Ethernet
400 Gigabit Ethernet
400G DCI
800G+ DCI
DWDM OTN
Wi-Fi Access Point
StarLight
CenturyLink
zayo
Internet2
ESnet
LEARN
Telia Carrier
SC18
Dallas, TX
hpc inspires.
General Conference Support
Chicago, IL
Akard St. POP, Dallas
DATABANK
noc-rtr-1
MX10008
140.221.239.253/32
2001:0468:1107:ffff::0253/128
noc-rtr-2
NCS5508
140.221.239.252/32
2001:0468:1107:ffff::0252/128
conf-rtr
MX960
140.221.239.255/32
2001:0468:1107:ffff::0255/128
SCinet NOC 2450
CalTech 1413
StarLight 2851
SCinet DNOC 420
SCinet DNOC 1034
SCinet DNOC 2316
SCinet DNOC 2644
SCinet DNOC 4026
USC Booth 1403
DevOps VMs
Reservoir Labs
REDSEAL
TACC
Cameleon
Gravwell
KAY BAILEY HUTCHISON
KBHCCD D326 IDF