



SCINET - THE WORLD'S FASTEST TEMPORARY NETWORK!

Networked Systems Tech Talk Series @ Illinois Tech

Angie Asmus, SC24 SCinet Chair, <u>angie.asmus@colostate.edu</u> Jason Zurawski, SC24 SCinet Executive Advisor, <u>zurawski@es.net</u>



AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- Digging Deeper: SCient Technology and SC23 Summary
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions



THE INTERNATIONAL CONFERENCE FOR HIGH PERFORMANCE COMPUTING, NETWORKING, STORAGE, AND ANALYSIS - SC23/SC24



- Long history, dating back to 1988
- Held annually in November
- Sponsored by the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE)
- Forum for researchers, engineers, and scientists to present their latest work, exchange ideas, and learn about the latest advances in high-performance computing and related fields.
 - Wide range of technical papers, tutorials, workshops, and panel discussions, as well as an exhibition hall where vendors showcase their latest products and services.
 - The exhibition hall also features cutting edge experiments and demonstrations powered by SCinet, the world's fastest temporary network.

SC24 SCINET CHAIR: ANGIE ASMUS





THANK YOU FOR A GREAT SCINET AT SC23 HANS ADDLEMAN!







AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- **Digging Deeper: SCient Technology and SC23 Summary**
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions





o ~14,000 Attendees SC22 - 11,830 Attendees SC21 - 6,694 (onsite in St Louis, hybrid)
 SC20 - 7,446 virtual SC19 - 13,918 (Denver)
 SC18 - 13,029 (Dallas)
 350+ exhibitors 200+ Technical Sessions, Workshops, and Birds of a Feather sessions 750+ total Volunteers SC23 wide



SC23 SCINET BY THE NUMBERS

- 6.71 Tbps delivered to the CCC
 16 400 Gbps connections
- 28.7kW power draw
- 206 Scinet Volunteers
- \$40.5m in donated & loaned hardware, software, and optics from 30 contributors
- 12+ miles of Fiber deployed
- 421 wireless access points deployed

WHAT DID WE DO LAST YEAR?

- IPv6 and Option 108
- XNet 5G Wireless
- 400G ZR+
- NSO Orchestration and Automation with EVPN/VXLAN
- Replaced Rack 1 panels with MTP Cassettes facing the show floor
- Stood up new management net rack-sw
- Migrated most management and commodity nets behind conf-fw for the first time
- Brand new racks from Schneider Electric/APC and new PDU
- Switch to INext
- Tapped a significant portion of NRE traffic for the first time, setting a record.
- We are probably the only people to ever tap a 400G ZR+
- 2x400GE (800GE) optics between vendors

- Multi-vendor zero touch provisioning
- BHR using FlowSpec
- RPKI (Alerting, not blocking)
- FABRIC in SCinet
- SCinet Research Exposition
- Automated BHR with FlowSpec via Gravwell and Corero
- 400G Miami to Denver service provided and managed by a contributor (Lumen)
- Tight integration between patch team and interconnects to leverage
- the excellent data presentation from iNext leading to a relatively smooth patching experience.
- IPv6 supporting legacy services via dns64 and nat64
- InfluxDB and Grafana



AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- Digging Deeper: SCient Technology and SC23 Summary
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions





WHAT IS DIFFERENT ABOUT SC?

- It has an exhibition where suppliers, researchers, and funding agencies all come together
- It has contributors that compete to provide products and services, but work together to create a solution for SCinet
 - opportunities for interoperability exploration
- It has a great set of people working together to create this solution

WHAT IS SCINET?



2 distinct networks built by volunteers

- Production network providing all wired and wireless internet access for attendees and exhibitors
 - Also provides bandwidth for live streaming of many sessions and workshops
- Research network in support of the Networked Research Exhibition (NRE) and SCinet's Experimental network projects (XNET).



WHAT IS SCINET?

1 year to design, 1 month to build, 1 week to operate, 1 day to teardow

SC23 Statistics

- 32 Network Research Experiments (NRE) on the show floor
 206 volunteers
- - 9 countries and 31 states represented
- 113 different organizations from around the world
 Since 2015, 50 WINS participants from 46 organizations and 43 states
 30 contributors loaning \$40.5M in hardware, software, and services
 9000#'s of Loaned Hardware

- 6.71Tbps of wide area network capacity
 - Includes 16 x 400Gbps connections and 1 x 800Gbps
- 420+ Wireless Access points deployed
- WiFi 6E, with IPv6 deployed to WiFi clients
- 5G Experimentation
- 12.65+ miles of fiber optic cable installed
- 2178 fiber patches
- 409 Helpdesk Tickets
- 33kw of power required for sustained operation (average home consumes this over a day)



AA

3/29/2024

SCinet Overview





SCINET CHALLENGES

- All volunteer team
- Almost all of our hardware is donated
 - Supply chains!
 - Drained demo pools
 - Sometimes get beta hardware
- Extremely tight timeline.
- Network Security!
 - SCinet/SC a very big target
- Customer service and performance troubleshooting until end of show



SCINET CONTRIBUTORS

SC23 Contributing Hardware, Software, and Service Vendors

- Cisco
- Pier Group / Aruba / HPE
- Arista
- Ciena
- ESnet
- Internet2
- Juniper
- Nokia
- Schneider / APC / Critical Components
- EXFO

3/29/2024

- Cornelis
- Gravwell
- Lumen
- Palo Alto
- Verizon
- TACC
- Corelight
- FRGP
- InMon
- Zayo
- Huber+Suhner Polatis
- Sentinel One



- University of Warsaw
- Florida International University

 \mathbf{O}

O

Coreo

FABRIC



AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- Digging Deeper: SCient Technology and SC23 Summary
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions





WHAT DO 200+ PEOPLE DO?

- SCinet is made up of 18 different teams with diverse skills sets
 - Architecture / Interconnection
 - Communications
 - Contributor Relations and Logistics
 - DevOps
 - Wireless and Edge
 - Networked Research Exhibition (NRE)
 - Experimental Networks (XNET)
 - Fiber
 - Help Desk
 - Network Security
 - Power
 - Routing

• Volunteer Services (Students)

- WAN Transport
- Finance
- Physical Security
- Digital Experience



WOMEN IN IT NETWORKING AT SC (WINS)

- The Women in IT Networking at SC (WINS) is a multi-year program initially funded by the National Science Foundation (NSF). *
- Was developed for addressing the prevalent gender gap that exists in Information Technology (IT), particularly in the fields of networking engineering and high-performance computing (HPC).
- Collaborative project managed by:
 - Indiana University
 - University Corporation for Atmospheric Research (UCAR)
 - Energy Sciences Network (ESnet)
- It was originally introduced as a pilot program** in November 2015 at the SC15, and has supported volunteers for SC16 - SC24.



WINS PROGRAM GOALS



- Expand the skill sets, professional network and communications skills, including public speaking, for the participants.
- Provide technical training that is applicable and valuable to the participants' home institution.
- Increase the diversity in the SCinet volunteer pool.
 - Since SC20 preference has been given to applicants who are historically underrepresented in the Information Technology field.
- Raise awareness of the issue of fewer women in IT.
- Gather program information, including candidate metrics and SCinet surveys (mentors and participants), to better understand gender diversity in the network engineering community and use the information to develop a sustainable program.





WINS ACTIVITIES

- Enables 3-8 women per year to participate in SCinet, a unique, hands-on experience in building an operational and research network from the ground up for the annual Supercomputing Conference
- Selected by a committee of experts and rigorously reviewed and evaluated
- SCinet provides an ideal professional development opportunity for engineers and technologists looking for direct access to the most cutting-edge network hardware and software.
- SC great partner and very supportive of inclusivity and increasing diversity
- Gives the opportunity to work side by side with the world's leading network and software engineers, and the top network technology vendors.
- Professional development activities fund WINS participants to attend and present at conferences to raise awareness of SCinet, diversity issues in IT, and technical content; alumni calls; The Quilt MOR Leadership program.









AA

3/29/2024

SCinet Overview

WINS AT SC24



WINS Returnees in SCinet Leadership

- SCinet Chair Angie Asmus
- Technical Director Brenna Meade
- Research Director Debbie Fligor
- Edge Network Team Chair Jennifer Kim
- Fiber Team Chair Julie Locke
- Interconnect Team Deputy Chair -Suzanne Prentice
- Network Security Team Deputy Chair Britt Huff
- Routing Team Chair Kalina Dunn
- WINS Team Chair Karen Lopez



New SC24 WINS Awardees: 6 3/29/2024



AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- Digging Deeper: SCient Technology and SC23 Summary
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions



SC23 WAN FACTS!

- 6.71 Tb/s
 - 16 400G connections, 1 800Gbps (+ some 100s, and 10s still!)
- Contributors
 - Ciena
 - Cisco
 - Nokia
 - Arista
 - Exfo
 - Viavi
 - H+S Polatis
- Circuit Providers
 - R&E_
 - ESnet
 - FRGP
 - Internet2
 - Commercial:
 - Lumen
 - <u>V</u>erizon
 - Zayo





•





H+P POLATIS: OPTICAL SWITCH

- First used in SCinet in 2019 •
- Allows users to remotely reroute optical • signals to different fibers
 - Some models include Optical Power Monitoring (OPM)
- Game-changer when it comes to testing circuits
 - Users are able to remotely loop or move • a circuit from a router to a test set for certification/troubleshooting and back without physically repatching the fiber
 - Saves an incredible amount of time and maintains fiber cleanliness
 - For SC23 we are working on interfacing with the API to allow for potential automation and to reduce the chance of unintentional service impact



IPV6



SCinet-wide initiative for SC23 ... but why?

- About 20% of the SCinet volunteers this year work for federal institutions. Federal mandate:
 - Ensure all new networked Federal information systems are IPv6enabled at the time of deployment to ensure the IPv6 only requirement is met. It is the agency's strategic intent to phase out the use of IPv4 for all systems. FERC requires that all IPv4 systems be upgraded to IPv6 upon the next acquisition cycle.
- SCinet is a great place for testing large changes and interoperability

We focused on 2 main parts of the network for IPv6-only:

- Management
- Wireless

BUILDING BLOCKS



Supporting services for IPv6 only Wireless: NAT64 DNS64 SLAAC and/or DHCPv6

DHCP option 108

Supporting services for IPv6 only management network: NTP Syslog Monitoring: SNMP, streaming telemetry **Configuration** backup NSO/automation tools/remote configuration tools ZTP







RFC 8925 IPv6-only Preferred option for DHCPv4 (https://datatracker.ietf.org/doc/html/rfc8925)

Using DHCPv4 to turn off IPv4 stack on the client for a specific interface

For networks that contain a mix of both IPv6-only-capable hosts and IPv4requiring hosts and that utilize DHCPv4 for configuring the IPv4 network stack on hosts

Using DHCPv4 for signaling is that IPv4 will be disabled only if both the client and the server indicate IPv6-only capability Caveat (1): not supported on all Client OS. e.g. does not work on Windows 10 or 11.... yet Caveat (2): This is a transitional tool for moving to IPv6 only networks



ALL CLIENTS ARE NOT CREATED EQUAL



MULTICAST TRAFFIC OVER WIRELESS



Our main technical issue we ran into was multicast issues on wireless which specifically affected wireless performance and timely router advertisements (Why is it always multicast?)

- Client isolation
- Multicast filtering

mDNS protocol traffic (really all broadcast/multicast discovery traffic including bonjour, LLMNR, Dropbox LAN sync, SSDP)







NOTABLE CLIENT ISSUES EXPOSED BY DHCPS OPTION 108

- Misconfigured, centrally based host firewalls caused significant issues with proper IPv6 functionality (i.e. do not filter ICMP)
- Certain older and non-google (pixel) Android devices have non-trivial IPv6 limitations causing an inability or limit in connectivity to certain IPv6 networks
- Router Advertisements use on-link multicast by default, leveraging unicast router advertisements may alleviate some multicast related issues





3/29/2024

S

NETWORKED RESEARCH EXHIBITION (NRE)



Network researchers and professionals from government, education, research, and industry are invited to submit proposals for demonstrations and experiments at the SC Conference that display innovation in emerging network hardware, protocols, and advanced network-intensive scientific applications.

- Researchers request SCinet services to their booth
- Use for projects requiring high-throughput traffic or private IP space
- Often implemented with VLANs (802.1Q) and loop prevention mechanisms (eg. spanning tree)
- Researchers collaborate requiring same traffic at multiple booths or between booths
 3/29/2024 Scinet Overview



NRE AT SC23



- 32 accepted experiments/demonstrations in areas including
 Data mobility and data transfer as a service for terabit network scales
 Single stream 400Gbps server capabilities
 International network peering at terabit scales
 Global data movement examples
 R&D efforts for the Large Hadron Collider (LHC)
 Packet marking and flow labeling
 Software Defined Exchange (SDX) Multi-Services for Petascale Science
 Testbed integration (e.g. FABRIC-Chameleon)
 Fully-Automated Network Configuration Management
 Multi-site data streaming orchestration
 International P4 Experimental Networks
 End-to-End Network Services and Workflow Integration
 Resilient Distributed Processing and Reconfigurable Networks
 Distributed Experimental Pipelines over WAN for Online Data Analysis
 5G on the Showfloor
 Quantum Networks

 - Quantum Networks Exploring FDT, QUIC, BBRv2 and HTTP/3 protocols in High Latency WAN paths Uncompressed 8K video processing across the WAN Distributed Machine Learning over Heterogeneous Edge Computing Resources 150+ VLANs implemented across multiple networks Many NRE demos had some type of tie to international networks.

inMon SC23 WAN Stress Test

SCinet Overview

SC23 DATA TSUNAMI

- As the showfloor is coming to a close, the NRE team coordinates an event where all experiments run simultaneously to maximize the WAN bandwidth utilization.
- Peak capacity of over 6 Tb/s was used during the Data Tsunami event at SC23.





Copyright © 1999-2023 InMon Corp. ALL RIGHTS RESERVED



INDIS AT SC23



- Innovating the Network for Data Intensive Science Workshop
 - Started in SC14 in New Orleans
- SC23 saw a half day workshop, and a half day demonstration session provided by the SCinet INDIS team
- Papers:
 - Enhancing perfSONAR Measurement Capabilities Using P4 Programmable
 - Data Planes (University of South Carolina) Experimental Study of TCP Throughput Profiles and Dynamics Over Dedicated Connections (Oak Ridge National Laboratory) Elephants Sharing the Highway – Studying TCP Fairness in Large Transfers
 - Over High Throughput Links
 - Imtiaz Mahmud (Lawrence Berkeley National Laboratory, University of Southern California, Renaissance Computing Institute, and Oak Ridge National Laboratory)
 - Evaluation of SCION for User-Driven Path Control (Politecnico di Torino, University of Amsterdam, and SIDN Labs)
 - Throughout Optimization with a NUMA-Áware Runtime System for Efficient Scientific Data Streaming (University at Buffalo & Argonne National Laboratory)

INDIS AT SC23 (CONT.)



- Distinguished speaker session from Ohio State University on "AI-EDGE: NG Networks meet Distributed Intelligence"
- Lightening Talks:
 - Testbed Evaluation of an Attestation-Capable, Programmable Software Switch (Illinois Institute of Technology)
 - Real-Time Frequency Moment Estimation on FPGA: Applications in Anomaly Detection and Weibull Flow Length Parameterization (Chung-Yuan Christian University)
 - Experimenting TCP Performance with Fabric (Lawrence Berkeley National Laboratory, University of Southern California, Renaissance Computing Institute, and Oak Ridge National Laboratory)

INDIS AT SC23 (CONT.)



- Demonstrations:
 - TCP experiment on FABRIC (Lawrence Berkeley National Laboratory, University of Southern California, Renaissance Computing Institute, and Oak Ridge National Laboratory)
 - DYNAMOS: Dynamically Adaptive Microservice-based OS A Middleware for Data Exchange Systems on FABRIC, (UvA)
 - FABRIC, ESnet, and NRP. ESnet SmartNIC based on U280 FPGA, then figured out how to also make this work on our U55C in NRP (IIT Chicago)
 - Towards Accountable Network Bandwidth Utilization via SDN (UCSD)
 - Testbed Evaluation of an Attestation-Capable, Programmable Software Switch: What is 'my' network equipment configured to do?, Nishanth Shyamkumar (Illinois Institute of Technology)



AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- Digging Deeper: SCient Technology and SC23 Summary
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions





4 PHASES OF SCINET (ON SITE)

Staging
Setup

NOC Move

Show
Teardown





SCINET STAGING



- Receive, inventory, and unpack contributor hardware, racks, power, etc
- Rack, stack, and power
- Install hardware, test, and configure at remote points of presence in show city.
- Configure, patch, test, test, test.









SCINET SETUP WEEK



- Prep for move to SCinet NOC on Exhibit floor.
- Move and lift racks on to the stage. (Wednesday)
- Power, repatch, and continue testing and turnup.
- Researchers start setting up toward end of week
- Go live is Saturday (before show) at noon
 - Conference starts Sunday morning 8am
 - Exhibit floor opens Monday night during gala







SCINET NOC MOVE CHALLENGES AND ADAPTABILITY



Staging area not always convenient to show floor Our most dangerous / intricate process. Takes nearly 2 days.

 must be careful with our volunteers and contributors equipment

3/29/2024

SCinet Overview

JZ



SCINET SHOW WEEK

- Managing, Monitoring, troubleshooting 2 networks for 5 days.
 - Monitoring and working with Digital Experience team to make sure our virtual attendees have a flawless experience (new in 21)
- Securing the network
 - Actively searching for threats internally and externally
- Help Desk for 1 on 1 customer support
 - Fielding tickets till very end of show
- NRE team works with each NRE participant to make sure their demonstrations are performant
- Adapting, changing, and adding as needed to keep up with capacity needs.



SCINET TEARDOWN

- Show Floor closes ~3pm Thursday
- NOC emptied out and all crates and boxes are moved into the space
- Show Floor Friday 8am
 - Factory reset equipment
 - Remove all patches
 - de-energize, remove from rack, inventory, pack up, ship.
- Commodity network must remain live till conference closes at Noon.
 - SC23 a small portion must remain live until 3pm for digital experience
- Friday 6pm Wrap dinner!





SCinet Overview



AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- Digging Deeper: SCient Technology and SC23 Summary
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions



WHAT'S TO COME FOR SC24!





49



SC24....ATLANTA, GA







Phil Roth Oak Ridge National Laboratory

SCinet Overview

AA



A LITTLE ABOUT GWCC

- Unique spaces
- Over 4m sq ft in exhibit and meeting space
- Separate Exhibit Halls (1m Sq Feet)
 SC24 will be in A Hall and B Hall

https://www.gwcca.org





5



VICTORIES, CHALLENGES, & INFO

- Fiber availability in and out of the building
- Shortened Staging Period
- Large distributed facility
- Challenges for wireless lots of area to cover, few copper connections to uplink to

SC24 INITIATIVES



- 1. Make Atlanta successful
 - Wired and Wireless Connectivity
 - WAN Connectivity
- 2. IPv6 continuation
- 3. SCinet Cluster Pilot
- 4. SCinet Education
 - Make sure SCinet has value for the volunteers
 - SCinet Theater Presentations, Communications to Organizations, Demonstrations & Talks
- 5. Diversity Initiatives (WINS, COMMUNITY)



AGENDA

- What is it?
- Fast Statistics
- Why SCinet?
- WINS
- Digging Deeper: SCient Technology and SC23 Summary
- SCinet: Year at a Glance
- SC24 Preparation
- Conclusions





FINAL THOUGHTS

- 1 year to design, 1 month to build, 1 week to operate, 1 day to teardown
- 2 distinct networks built by volunteers
 - Production network
 - Research network
- 200+ Volunteers, from 100+ organizations across 18 teams working together to design, build, and operate the world's fastest network
- Focus on education, diversity, and innovation







SCINET - THE WORLD'S FASTEST TEMPORARY NETWORK!

Networked Systems Tech Talk Series @ Illinois Tech

Angie Asmus, SC24 SCinet Chair, <u>angie.asmus@colostate.edu</u> Jason Zurawski, SC24 SCinet Executive Advisor, <u>zurawski@es.ně</u>t