



SC20

Everywhere | more
we are | than hpc.

Future of CI and role of XNET
Cees de Laat – Univeristy of Amsterdam

2020-11-13 • virtual (everywhere)

History of XNET innovations

SCinet's Xnet (eXperimental Networks) environment provides a venue to showcase emerging, often pre-commercial or pre-competitive, developmental networking technologies, protocols and experimental networking applications.

SCinet is the Exhibit floor network and it has evolved into a robust, high-performance, production-quality network that exhibitors and attendees depend on for reliable local area, and wide area network services. Consequently, it has become increasingly difficult for SCinet to showcase cutting edge, potentially fragile technology. Xnet provides a venue which is by definition cutting-edge, pre-standard and in which fragility is understood and accepted. Xnet enables its participants and research exhibitors an opportunity to showcase emerging, prototype or experimental network gear or capabilities, prior to their general commercial availability in a fault tolerant (forgiving) environment. Participants in Xnet's environment gain insight from the worlds most advanced computing and computer networking audience.

Xnet debuted in Portland, Oregon at SC99, where **Dense Wavelength Division Multiplexing (DWDM)** technology was used in the implementation of OC-48 SONET rings on the conference show floor. At SC2000, Xnet demonstrated pre-production and early delivery **10-Gigabit Ethernet** equipment connecting several exhibit floor booths. Since then, Xnet has showcased many advanced optical switching technologies as well as international, multi-domain control plane manipulation and new transport technologies. These have included 16 - 64k pMTU (SuperJumboFrames), 10G Encryption and Remote Direct Memory Access over InfiniBand in 2005 as well as GMPLS control, Sensor Networks and Data Centre o/s migration, and open fabric NFS/RDMA using RDMA over Infiniband

Apart from firewalling advanced technologies from the production SCinet, XNET has some specific advantages for participants. We offer show-floor booth space for demo's that have no home booth. There is significant attention to 'what's new' and XNET is one of the must-see go-to activities. We also include descriptions in show media both print and electronic. During SC, there are numerous directional and promotional signs that tell visitors what is in XNET and where the demos are located. From both a research and marketing perspective, XNET brings a great deal of attention to our participants.

sc06

- SCinet's Xnet (eXtreme net) provides a venue to showcase emerging, often pre-commercial or pre-competitive, developmental networking technologies, protocols, and experimental networking applications.
- Xnet is an opportunity to showcase emerging, prototype or experimental network gear or capabilities, prior to their general commercial availability in a fault tolerant (forgiving) environment. Participants in SCinet's Xnet environment gain insight from the worlds most advanced computing and computer networking audience.
- DRAGON Xnet demo
details: <http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/SuperComputingPlanning2006>
- GSFC's High End Computer Network (HECN) Team supported four realtime high performance networking data flow demonstrations into the showroom floor of the International Conference for High Performance Computing, Networking and Storage, a.k.a. SC2006, hosted in Tampa, FL. The provided support was in the form of either HECN's physical network infrastructure used in the critical path of a demo's realtime data flows or network engineering or troubleshooting expertise to help setup the demo.
- The demos supported were those of the following projects:
 - DRAGON: <http://dragon.maxgigapop.net/>
 - TeraFlow Testbed: <http://www.teraflowtestbed.net/>
 - OptIPuter: <http://www.optiputer.net/>
 - DICE: <http://www.avetec.org/dice/>

SPEED was the problem!

SC08 Xnet Proposed Demo description

High Performance Digital Media Network Demonstration

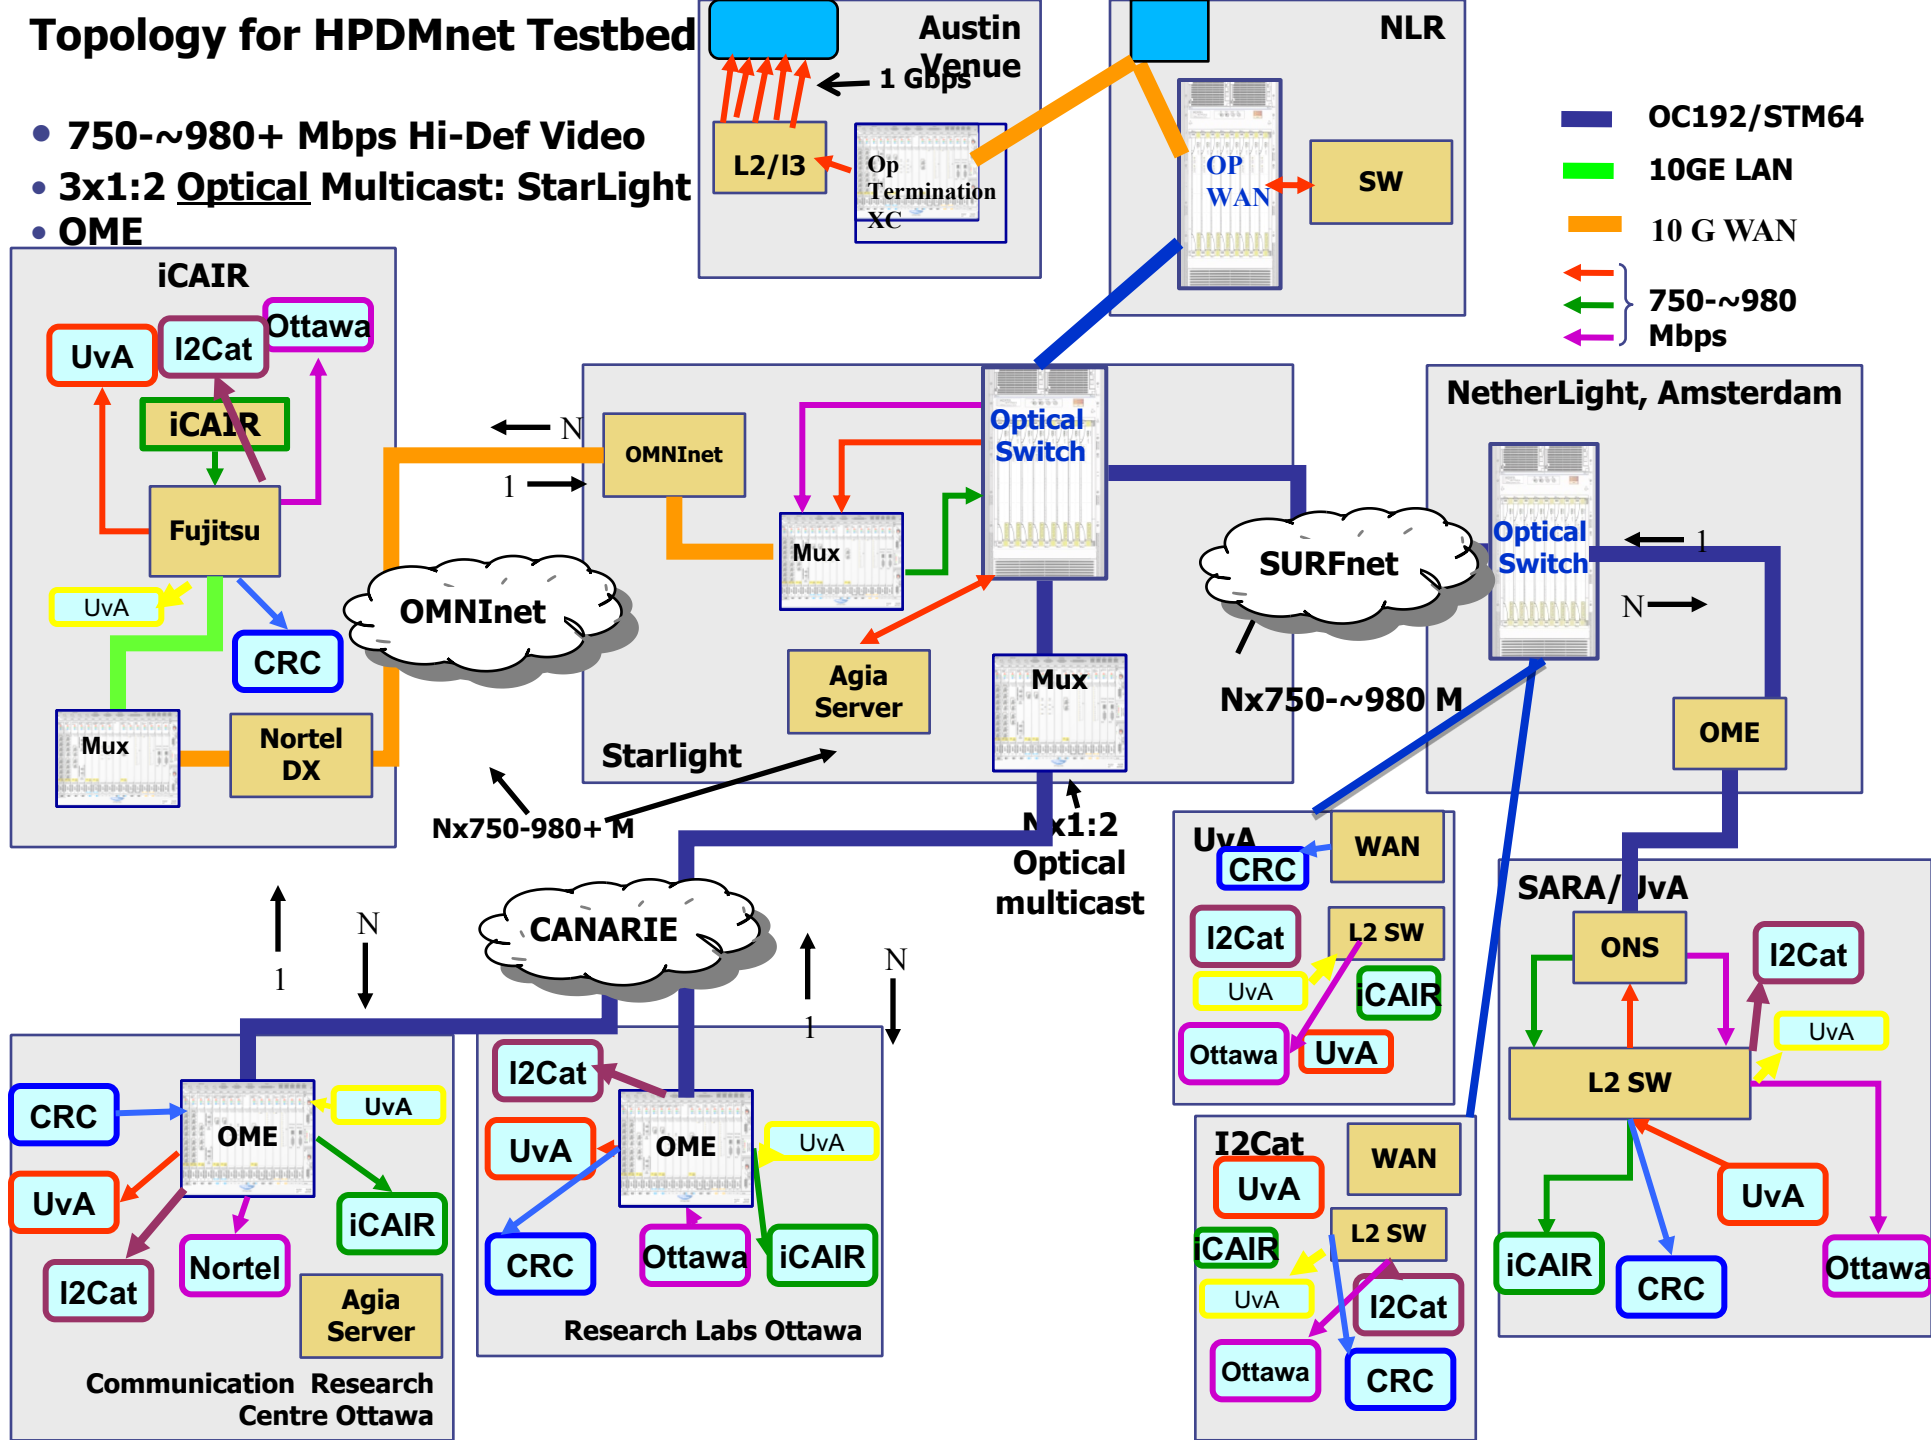
- The High Performance Multi-media project is a collaboration between government, academic and industrial researchers where this international advanced networking demonstration will showcase capabilities of the High Performance Digital Media Network (HPDMnet). For several years, the HPDMnet community has designed, developed, and demonstrated dynamic lightpath provisioning capabilities, for both inter and intra domain implementations. Now, they are developing advanced services that can build upon those capabilities, by optimally utilizing them. For example, the HPDMnet initiative is creating the world's first international high performance transport service specifically created for high quality, large-scale digital media, including support for extremely high volume media streams. This initiative represents a complete departure from existing services, which are based on legacy services and technologies and which cannot meet many emerging needs for high quality, high performance reliable services. This initiative is providing an infrastructure foundation for future digital media services, as well as for other data-intensive applications, which will be first implemented at Global Lambda Integrated Facility (GLIF) international exchange facilities (GOLES). To enhance such services, the HPDMnet community is participating in the development of an architectural framework for a high performance network interface, a GLIF Network Interface, which will be standardized through the OGF process developing the Grid Network Interface architecture. This initiative was established as a cooperative international partnership by several major network research centers. Members of the consortium include: CANARIE, Communications Research Centre (CRC) Canada, I2Cat, Inocybe, International Center for Advanced Internet Research, Korean Institute of science and Technology Information, Northwestern University, NetherLight, Nortel, SARA, StarLight, SURFnet, and University Van Amsterdam.

- The key point is the nature of the service that will be demonstrated, in part, because these capabilities are being designed and implemented as a real, defined "service," and not merely as a collection of technologies that are being showcased. The supporting technologies are certainly important - they provide the enabling resources. However, the service is based on innovative methods that provide for unique integration, control and management of lower level technology-based capabilities providing for capacity, performance, dynamic lightpath provisioning, stream duplication, and other functions. The modularity of the architecture allows for additional capabilities to be added with minimal or no manual intervention. Primarily, this is accomplished by using XaaS techniques (SaaS, PaaS, IaaS, etc), that can be highly decentralized. In addition, this service is being designed to be persistent at major exchange points, not merely a one-off demo. Also, in the final analysis, such services are judged by what they can do for applications. In this case, this service will be used to demonstrate high quality, high performance, data intensive digital media L1/L2 transport across national and international WANs, both live and streamed from repositories.

GRID / OGF was still big back then!

Topology for HPDMnet Testbed

- 750~980+ Mbps Hi-Def Video
- 3x1:2 Optical Multicast: StarLight
- OME



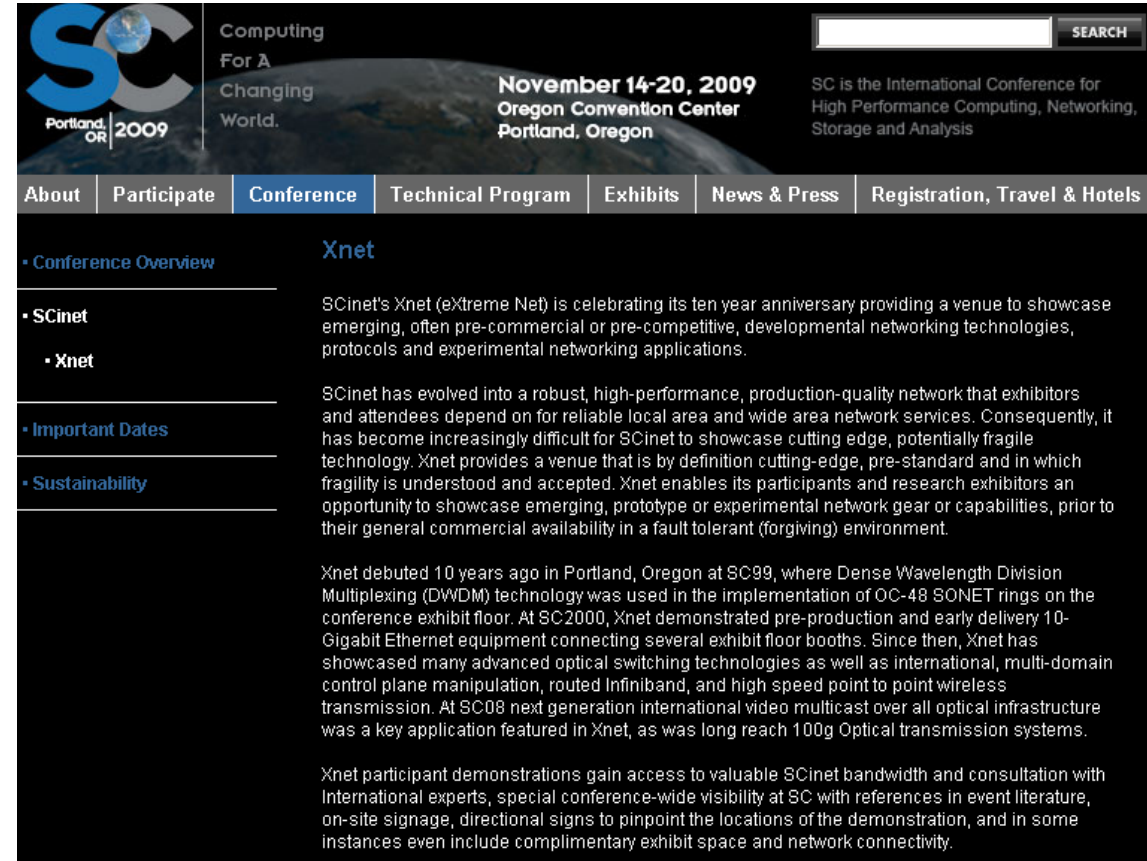
SC09

Scope:

- Experimental networking techniques & architectures
- Switching, provisioning, control
- Novel Applications, use of network
- Advanced Protocols
- Transmission, reach, speed

Result:

- nil
- Function of:
 - research shift to applications
 - reduced University research in “telecommunications”
 - Intense competition between Vendors
 - Vendor’s marketeers in charge (exhibition corrupts product launch)
 - Economic downturn



Computing For A Changing World. November 14-20, 2009 Oregon Convention Center Portland, Oregon

SC is the International Conference for High Performance Computing, Networking, Storage and Analysis

About Participate **Conference** Technical Program Exhibits News & Press Registration, Travel & Hotels

• Conference Overview

• SCinet

• Xnet

• Important Dates

• Sustainability

Xnet

SCinet's Xnet (eXtreme Net) is celebrating its ten year anniversary providing a venue to showcase emerging, often pre-commercial or pre-competitive, developmental networking technologies, protocols and experimental networking applications.

SCinet has evolved into a robust, high-performance, production-quality network that exhibitors and attendees depend on for reliable local area and wide area network services. Consequently, it has become increasingly difficult for SCinet to showcase cutting edge, potentially fragile technology. Xnet provides a venue that is by definition cutting-edge, pre-standard and in which fragility is understood and accepted. Xnet enables its participants and research exhibitors an opportunity to showcase emerging, prototype or experimental network gear or capabilities, prior to their general commercial availability in a fault tolerant (forgiving) environment.

Xnet debuted 10 years ago in Portland, Oregon at SC99, where Dense Wavelength Division Multiplexing (DWDM) technology was used in the implementation of OC-48 SONET rings on the conference exhibit floor. At SC2000, Xnet demonstrated pre-production and early delivery 10-Gigabit Ethernet equipment connecting several exhibit floor booths. Since then, Xnet has showcased many advanced optical switching technologies as well as international, multi-domain control plane manipulation, routed Infiniband, and high speed point to point wireless transmission. At SC08 next generation international video multicast over all optical infrastructure was a key application featured in Xnet, as was long reach 100g Optical transmission systems.

Xnet participant demonstrations gain access to valuable SCinet bandwidth and consultation with International experts, special conference-wide visibility at SC with references in event literature, on-site signage, directional signs to pinpoint the locations of the demonstration, and in some instances even include complimentary exhibit space and network connectivity.

Did not work out.

SC10 → Research Sandbox

- The SCinet Research Sandbox Initiative is a new feature of SCinet, created to allow researchers, computer scientists and network engineers to experiment using the state-of-the-art SCinet as their laboratory. The purpose of the SCinet Research SCandbox is to create new uses of SCinet from within the Research community at SC, to draw special attention to research demos that leverage SCinet in novel ways, and to better link SCinet to the Technical Program through Researcher's papers, presentations and posters that include descriptions of how SCinet helped.
- Interested participants must submit a proposal that describes their research project, SCinet network topology and description of how the research or results will be communicated, presented or reviewed. Special attention must drawn to how SCinet enhanced, accelerated or improved the research and demonstration. This description is the primary deciding factor in making a research proposal part of the sandbox program.
- [SCinet Sandbox August update.ppt](#)
- [Network Security Analysis proposal.pdf](#)
- [DICE 2010 Sandbox Proposal v0_3.pdf](#)
- [Caltech 40GE-Remote.png](#)



Connecting communities through HPC

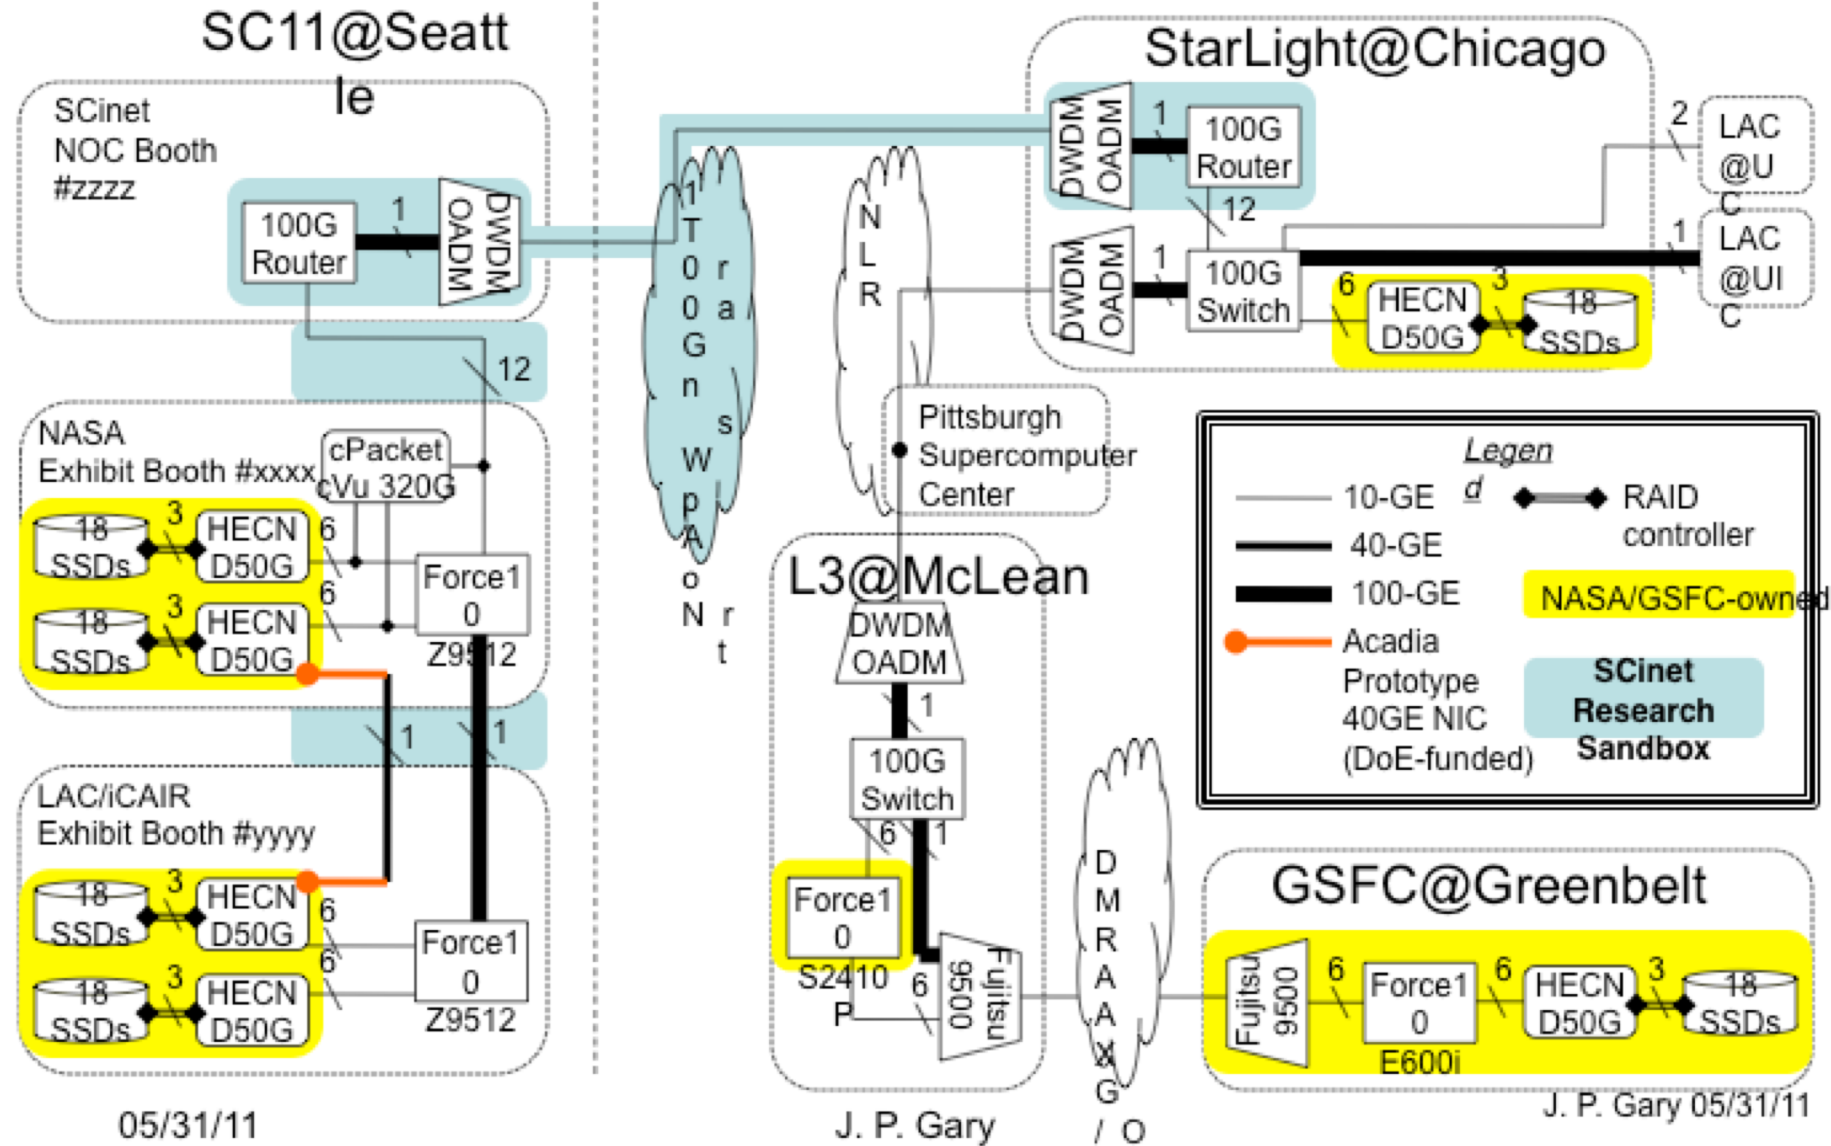
- Received 13 submissions
 - 2 on 100G
 - 1 on 40G RDMA over the WAN
 - 1 on passive monitoring
 - 9 on OpenFlow
- Accepted 11
 - Top rated 6 will get a chance to speak in the technical program



Evaluations/Demonstrations of 40-to-100 Gbps IPv4/IPv6 Disk-to-Disk File Copying Performance Across LANs & WANs

A Part of the 100G Consortium Initiative Among NASA and Several Partners

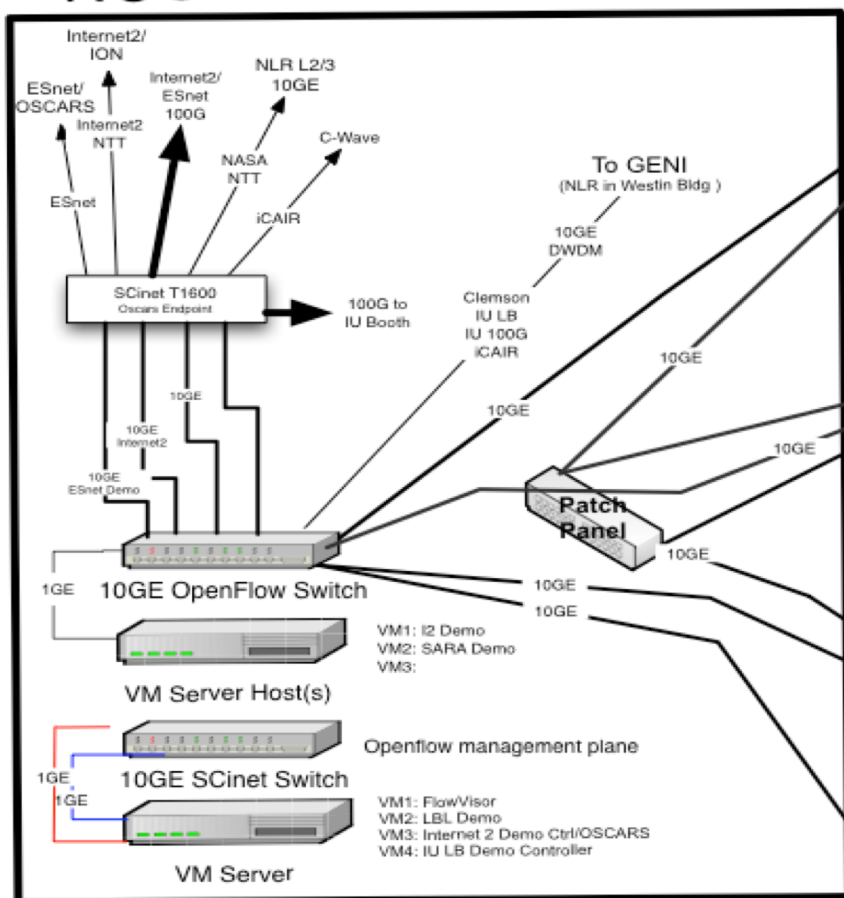
SC11@Seatt



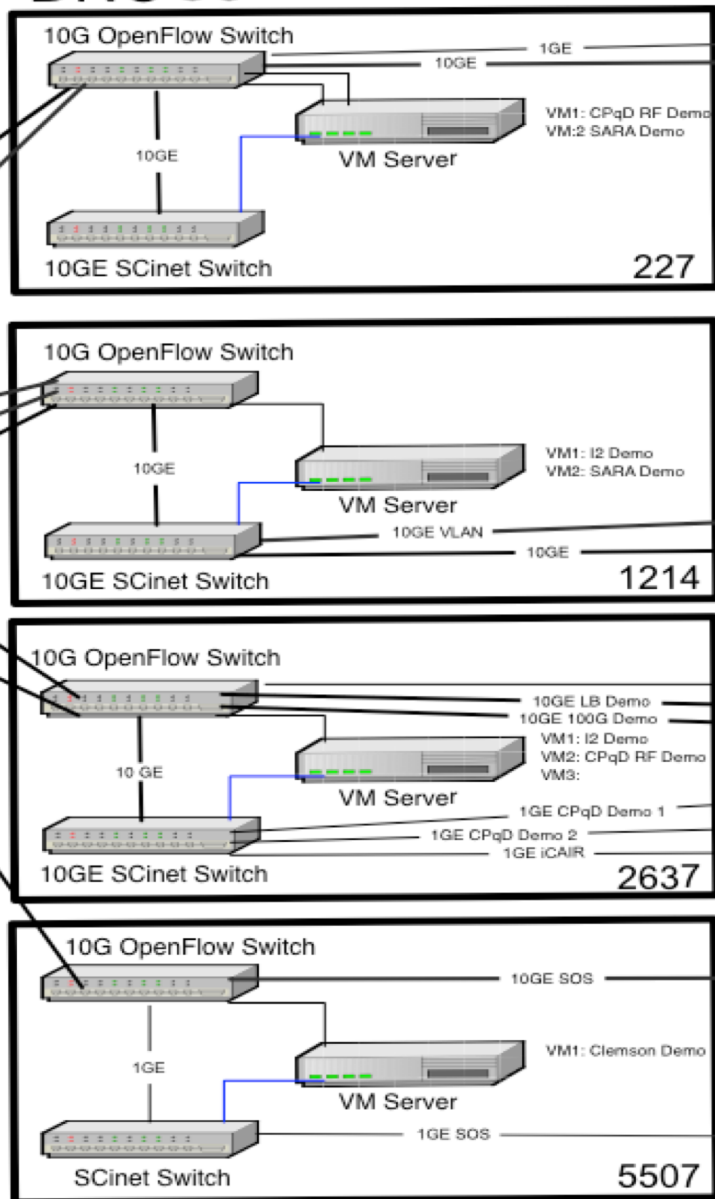
SCinet Research Sandbox OpenFlow Testbed (SC11)

Booth Openflow connection (need to confirm):
 - VLAN on booth connection: LBL, Internet2, DRC, iCAIR

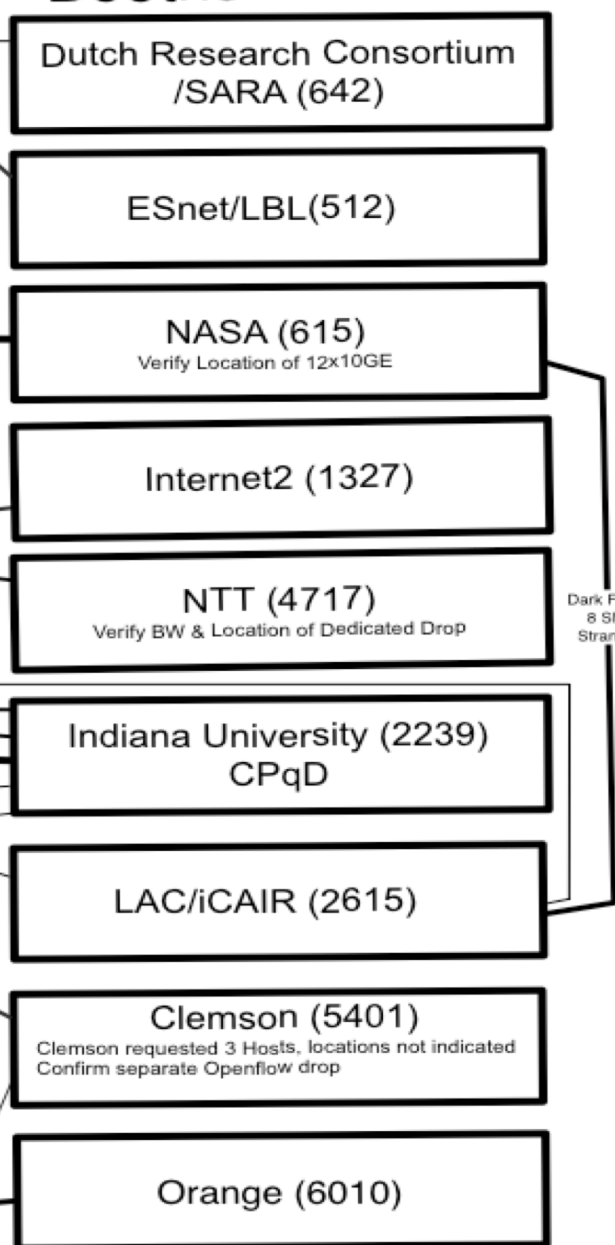
NOC



DNOCs




Booths



Dark Fiber
12 x 10GE

Dark Fiber
8 SM Strands

OpenFlow Management Network



Sc14 ==> INDIS on board + NRE – no specific XNET

- Paper centric
 - Cees DeLaat
 - Jennifer Schopf
 - Martin Swany
- Demo centric
 - Brian Tierney
 - Matt Zekauskas
 - Cees DeLaat
 - And a host of SCinet committee members to review
- Generally Helpful
 - Mary Hester
 - Kate Mace
 - Laura Diaz

Demo panel discussions will include:

- Optical networking
- Software-defined networking
- Network monitoring and traffic analytics
- Science DMZs and other campus network constructs
- Application-based data movement
- Open clouds and storage area networks

Paper panel discussions will include:

- OpenFlow and Software Defined Networking
- Use of alternative protocols and large-scale data movement approaches
- Advanced tools to support faster campus data networks
- Tuning networks for specific use cases, such as visualization
- Adaptations of OpenFlow to work within cloud environments

SC15-18, SC15-17 were INDIS-NRE efforts, no specific XNET, however, SC18

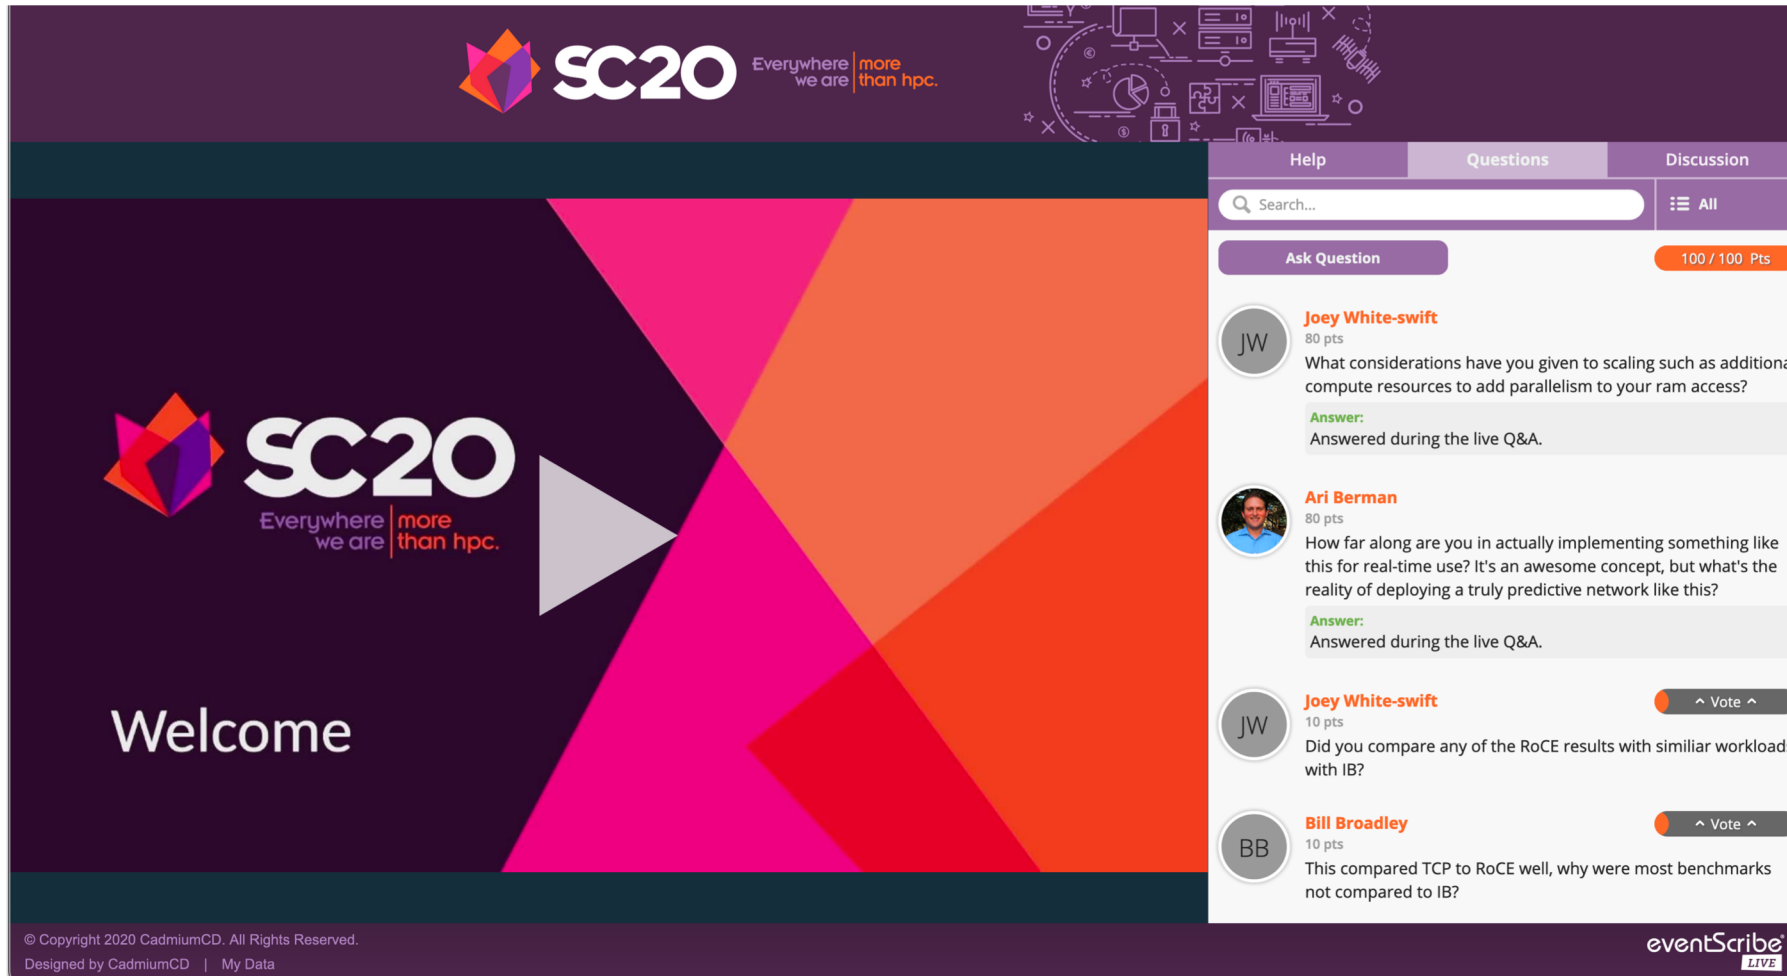
However, SC18:

- DTN Team
- Faucet Team
- Optical Cross Connect Team
- 400GE
- Controllers of Controllers (This experiment is hibernating this year.)

- Intro of Technical Challenge
- SC19 Scinet Data Transfer Node (DTN) as a service
- Streaming Telemetry for network performance visualization and alerting
- Optical Cross-connects
- SCION (Secure multi-path routing)
- Netbox
- Global Virtualization Services (GVS) - Network Virtualization

SC20 ==> Conference Virtualisation

Total SCinet Virtualisation! ;-)




The screenshot displays the SC20 virtual conference interface. The top header features the SC20 logo and the tagline "Everywhere we are | more than hpc." The main content area is split into two sections. On the left, a large graphic with the SC20 logo and tagline is accompanied by a large grey play button icon and the word "Welcome". On the right, a navigation menu includes "Help", "Questions", and "Discussion". Below the menu is a search bar and an "Ask Question" button with a "100 / 100 Pts" indicator. The "Questions" section lists three questions with their respective authors and points:

- Joey White-swift** (80 pts): "What considerations have you given to scaling such as additional compute resources to add parallelism to your ram access?"
Answer: Answered during the live Q&A.
- Ari Berman** (80 pts): "How far along are you in actually implementing something like this for real-time use? It's an awesome concept, but what's the reality of deploying a truly predictive network like this?"
Answer: Answered during the live Q&A.
- Joey White-swift** (10 pts): "Did you compare any of the RoCE results with similiar workloads with IB?" (Vote button)
- Bill Broadley** (10 pts): "This compared TCP to RoCE well, why were most benchmarks not compared to IB?" (Vote button)

At the bottom, the footer contains copyright information: "© Copyright 2020 CadmiumCD. All Rights Reserved. Designed by CadmiumCD | My Data" and the eventScribe LIVE logo.

Title: Future of CI and role of XNET - SCinet in the Roaring Twenties...

- Well, the Twenties 100 years ago ended badly, this century has a challenging start
- The aim of this presentation is to think about how cyber infrastructure to support science will evolve in the twenties and what SCinet may consider as platform offerings supporting booth demonstrators in 2025 and 2030.
- Before predicting the future, let us look in transformational developments in
 - Computing
 - Networking
 - Data
 - Science



Some progress



Apple Watch 4
On my wrist

2018

= ~7x

~ 540 MHz
? MFlops
1000 MByte memory
16000 MByte ssd
0,0012 kWh – 18 h



1976
Los
Alamos

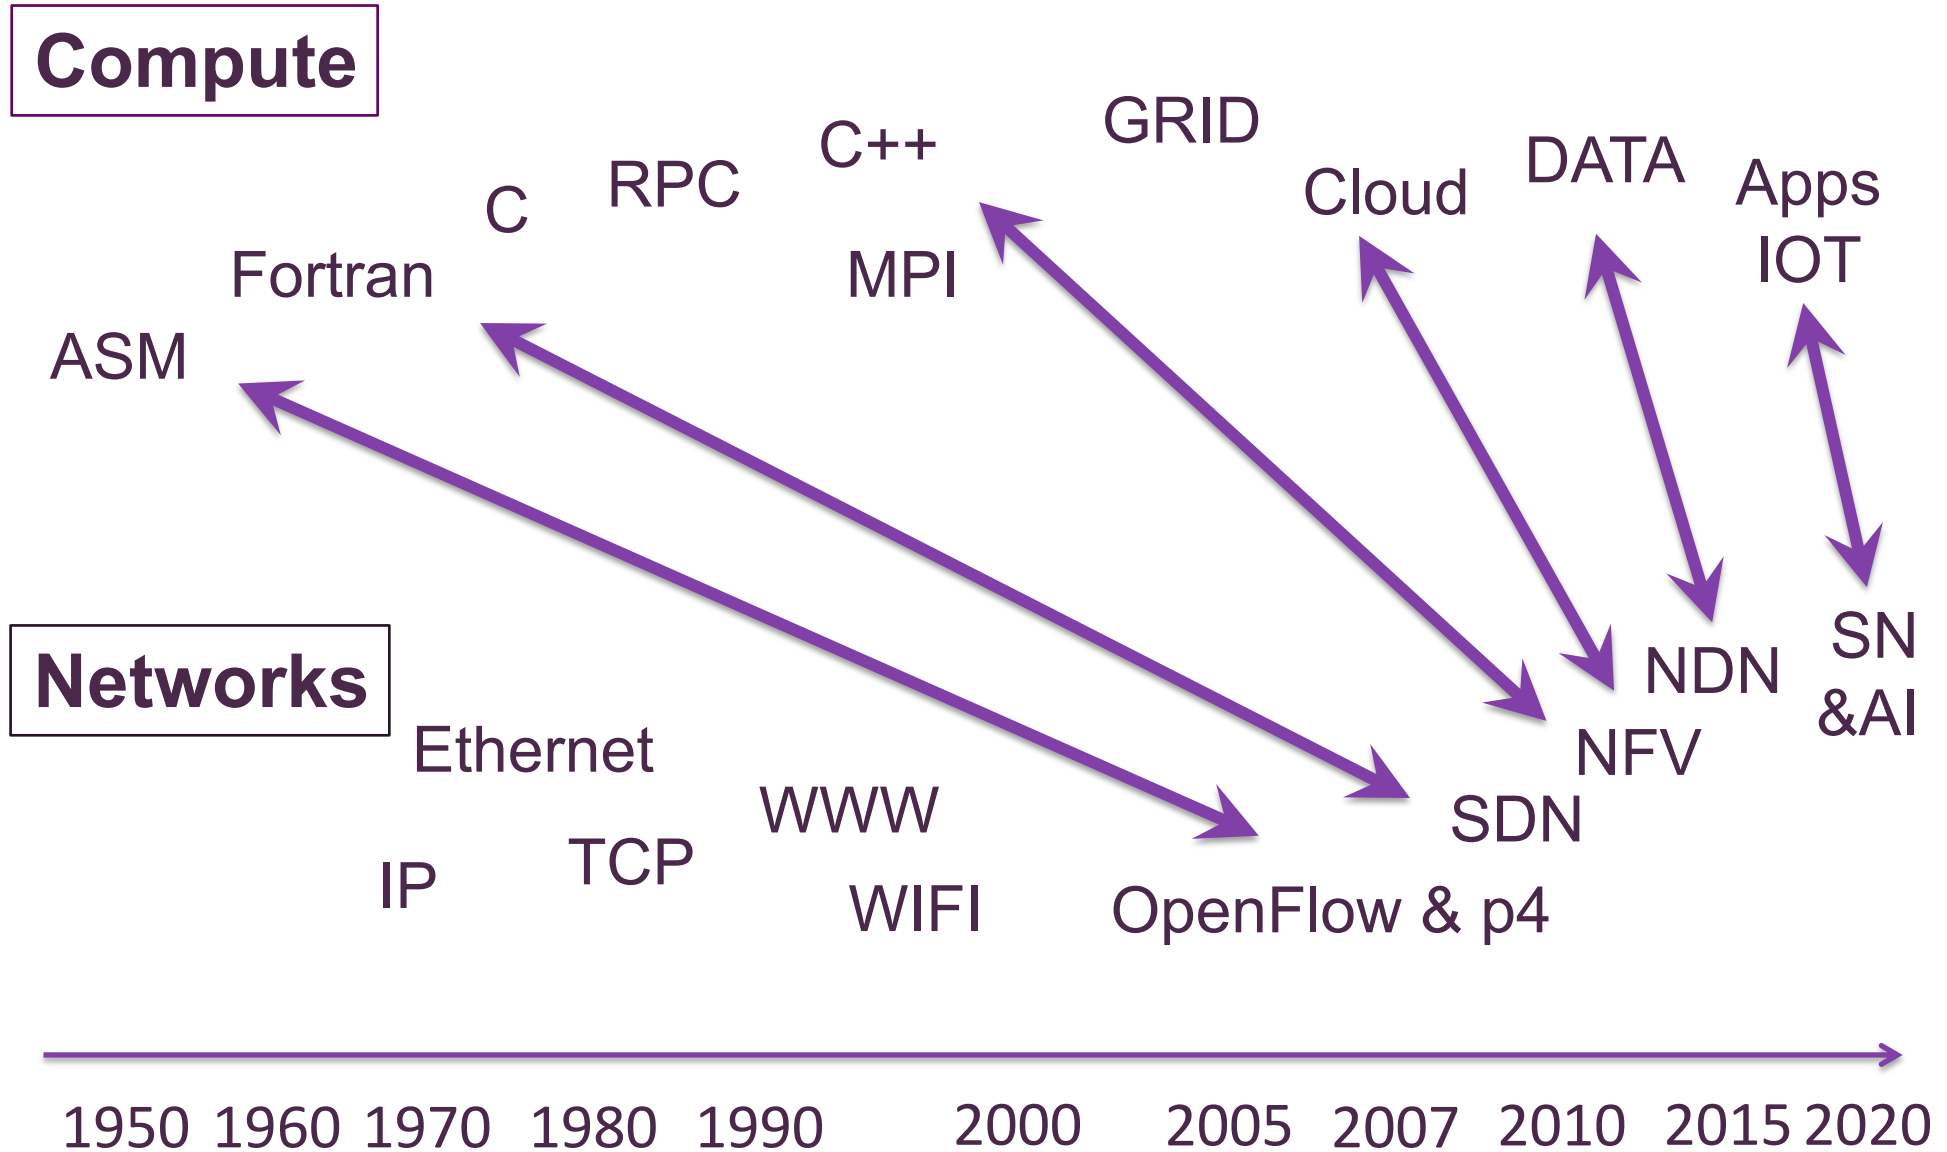


80 MHz
160 MFlops
8 MByte memory
300 MByte disks
120 kW

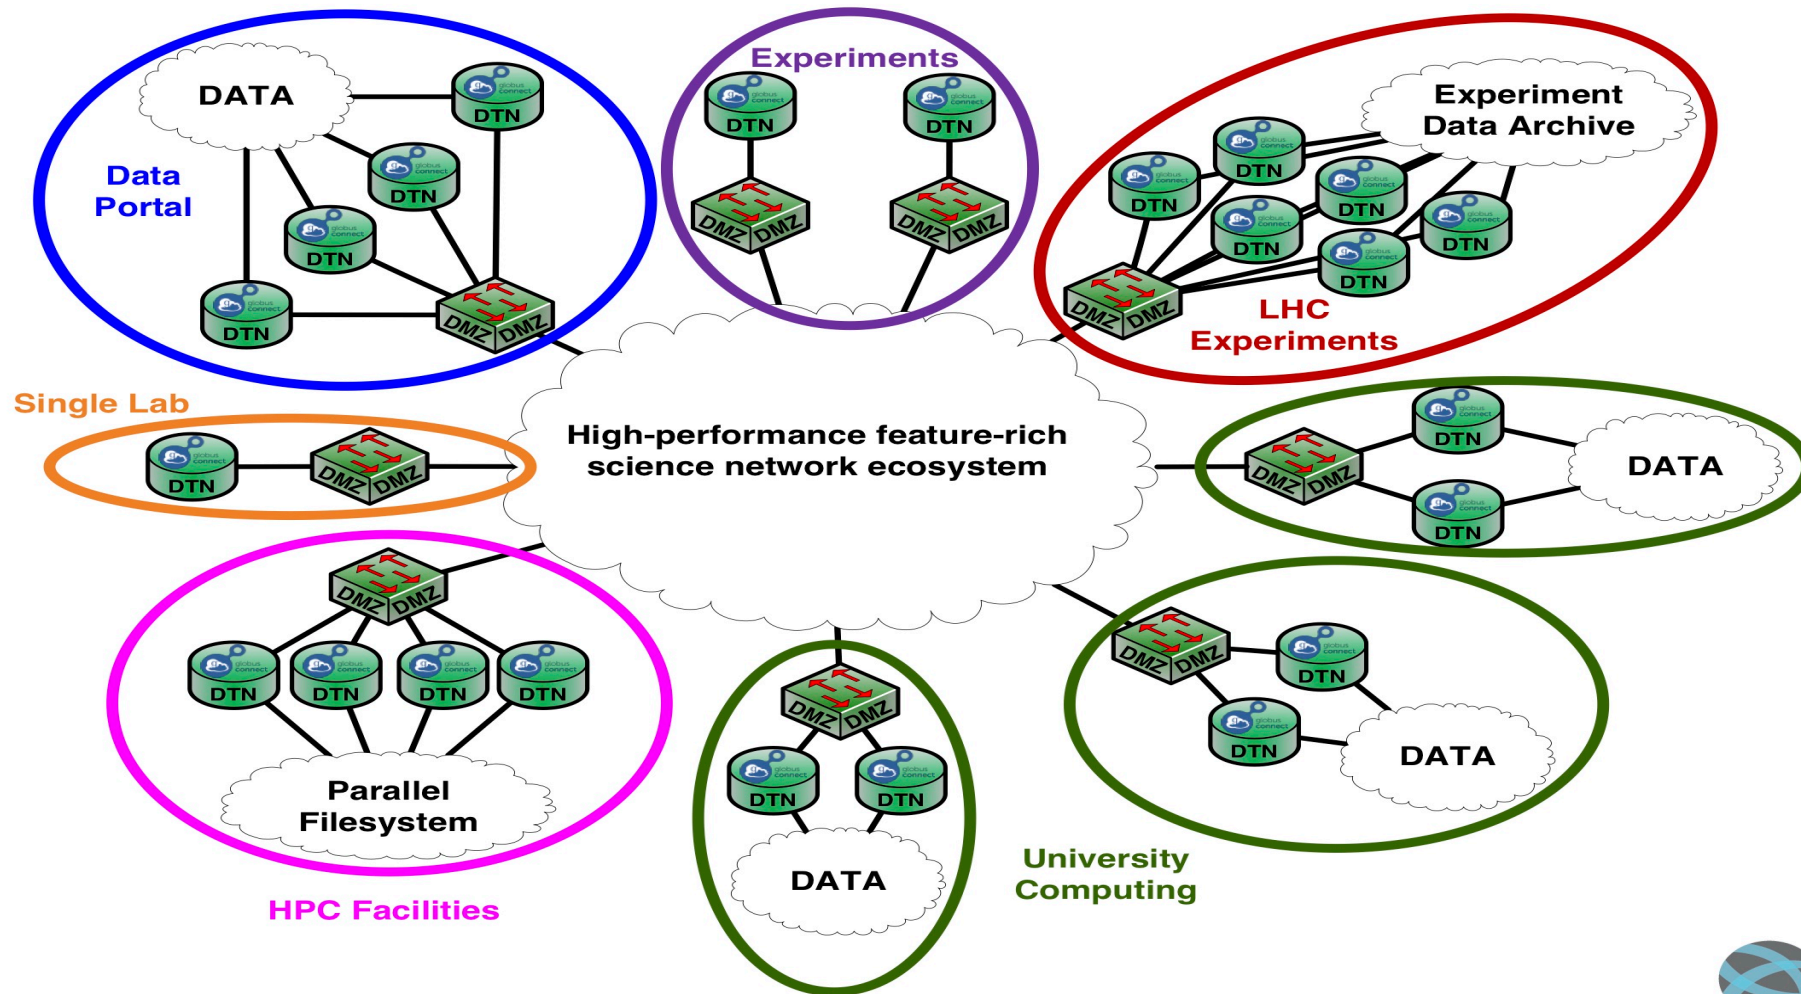
IOT & BYOD



Timeline

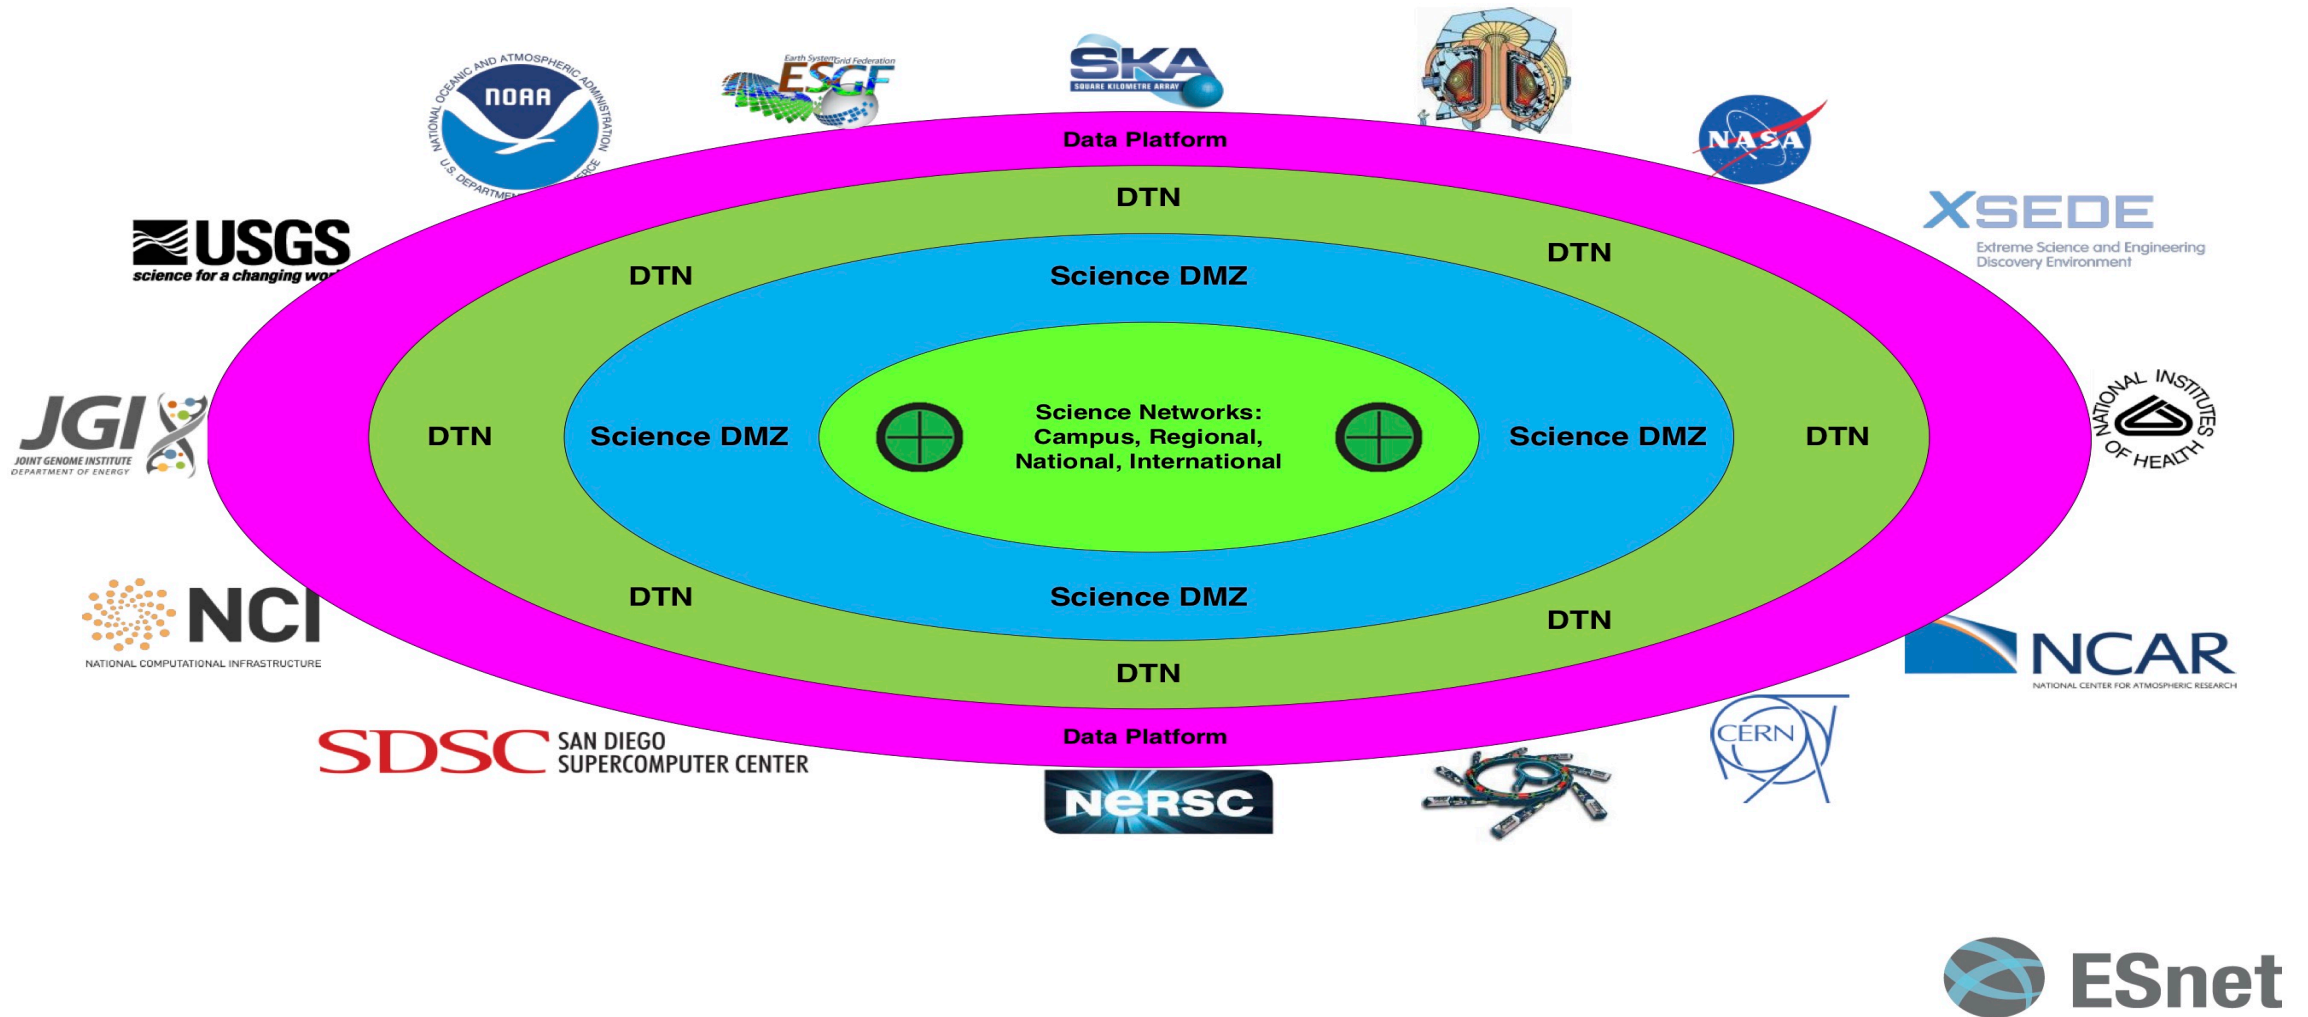


Science DMZs for Science Applications



Data Ecosystem – Concentric View

Looks a lot like an internet of things, but very big things! 😊



Responsible and transparent secure data layer



AMDEX

Data objects & methods
Data & Algorithms service

FAIR / USE


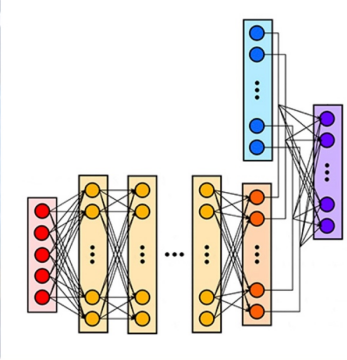

AMS-IX

Routers - Internet – ISP's - Cloud
IP packet service

IP / BGP

Layer 2 exchange service
Ethernet frames

ETH / ST



Scientific Machine Learning & Artificial Intelligence

Scientific progress will be driven by






- Massive data: sensors, simulations, networks
- Predictive models and adaptive algorithms
- Heterogeneous high-performance computing

Trend: Human-AI collaborations will transform the way science is done.



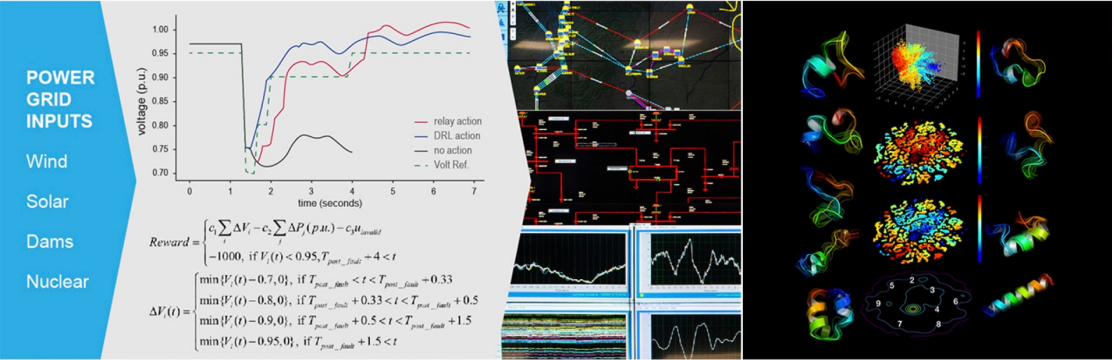
BASIC RESEARCH NEEDS FOR Scientific Machine Learning

Core Technologies for Artificial Intelligence

EXEMPLARS OF SCIENTIFIC ACHIEVEMENT



Human-AI insights enabled via scientific method, experimentation, & AI reinforcement learning.



Prepared for U.S. Department of Energy Advanced Scientific Computing Research




DOE Applied Mathematics Research Program
Scientific Machine Learning Workshop (January 2018)

Workshop report:
<https://www.osti.gov/biblio/1478744>

Title: Future of CI and role of XNET - SCinet in the Roaring Twenties...

- Well, the Twenties 100 years ago ended badly, this century has a challenging start
- The aim of this presentation is to think about how cyber infrastructure to support science will evolve in the twenties and what SCinet may consider as platform offerings supporting booth demonstrators in 2025 and 2030.
- Before predicting the future, let us look in transformational developments in
 - **Computing → From Dinosaurs to Ant Colonies**
 - **Networking → From end to end networking to Data Center centric**
 - **Data → From putting Data in the Cloud to Peer to Peer**
 - **Science → From IOT-Instrument via Fabric & AI to FAIR and Publication**
- **SO WHAT?**



SC2x

The way forward:

- Towards full programmability & virtualisation of networks
- Many booth demonstrators are on
 - computational steering of models with data
 - Stream processing
 - Artificial Intelligence & Machine Learning
 - Data distribution and Storage
 - Integration of IOT with data fabrics
- Very few are on speed of transfer & parallelisation of flows
- Introduction of Data Aware Layer
- Introduction of DTN and AI facilities in the network fabric
- Security, transparency and integrity major concerns in the 20's
- Complexity and control → self driving ~~cars~~ facilities

SCinet & XNET in 5 to 10 years


- Technical and capability challenges
- Introduction of DTN as a Service
- Harbour for Docker Containers and VM's
- AI & ML on demand (maybe, maybe not)
- Security functions

SCinet adopting select data functions, that is the question?

Questions? In the Q & A and Panel.

- This contribution was supported by the Dutch Science Foundation in the Commit2Data program projects dl4ld (grant no: 628.001.001) and EPI (grant no: 628.011.028).

Thanks for your attention!



See: delaat.net/dl4ld , delaat.net/epi and sc.delaat.net

