Networks for HEP and Data Intensive Science and the Digital Divide

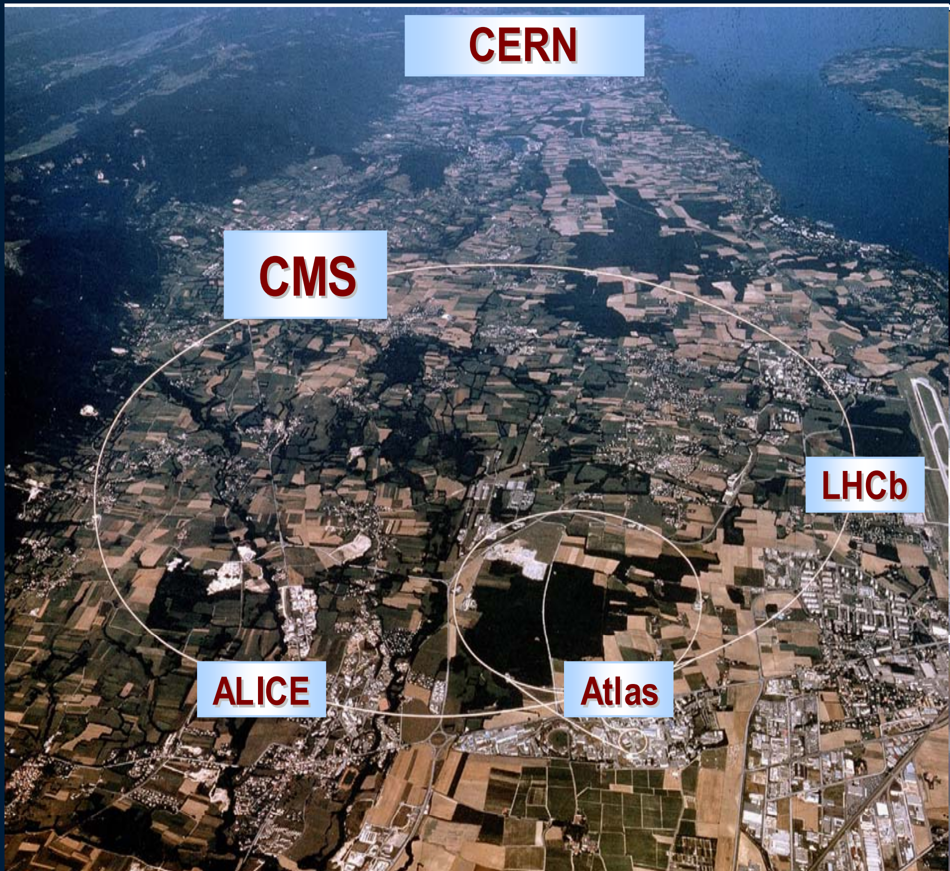

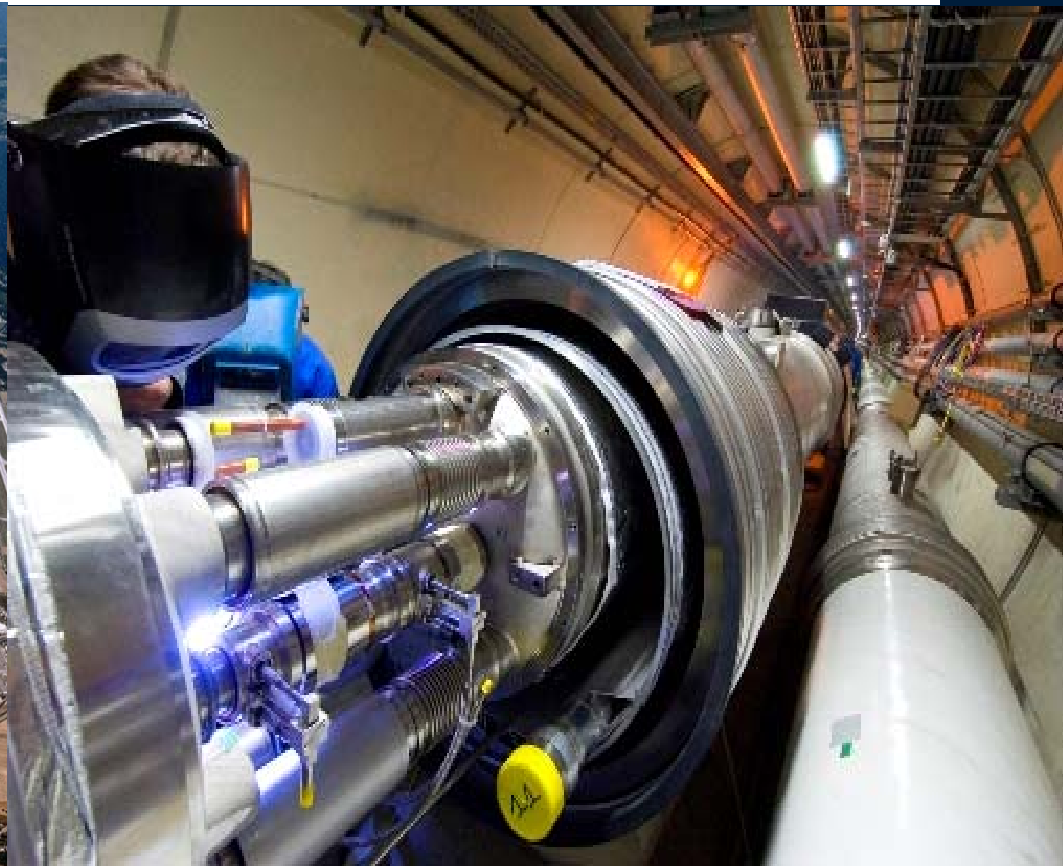

Harvey B Newman, Caltech

CHEP07

Victoria, September 5, 2007

ICFA SCIC

ICFA Standing Committee on Interregional Connectivity (SCIC)

- 1996 Visionary Statement of ICFA: Our Major Collaborations Should Gear Themselves for Remote Participation
- 1997 Network Task Force [Dedicated to David O. Williams]
- SCIC Created in July 1998 in VancouverMake recommendations to ICFA concerning the connectivity between the Americas, Asia and Europe
- As part of the process of developing these recommendations, the committee should
  - Monitor traffic on the world's networks
  - Keep track of technology developments
  - Periodically review forecasts of future bandwidth needs, and
  - Provide early warning of problems

SCIC in 2006-2007
http://cern.ch/icfa-scic

Three 2007 Reports: An Intensive Year
Rapid Progress, Deepening Digital Divide

- ***Main Report: "Networking for HENP"*** **[H. Newman, et al.]**
  - ➔ **Includes Updates on the Digital Divide, World Network Status; Brief updates on Monitoring and Advanced Technologies**
  - ➔ ***31*** **Appendices: A World Network Overview** ***Status and Plans for the Next Few Years of Nat'l & Regional Networks, HEP Labs, & Optical Net Initiatives***
- ***Monitoring Working Group Report*** **[L. Cottrell]**

Also See:

- **TERENA (www.terena.nl)** ***2005 and 2006 Compendiums:*** **In-depth Annual Survey on R&E Networks in Europe**
- **http://internetworldstats.com: Worldwide Internet Use**
- ***SCIC 2003 Digital Divide Report*** **[A. Santoro et al.]**

1st Revolution: "Long Dawn" of the Information Age

1.17B Internet Users; 300M with Broadband (7/07)

http://internetworldstats.com

- Explosion of bandwidth use: ~1 TByte/sec
- Raw capacity still largely unused
- Emergence of Web 2.0: Billions of Web Pages, rich content, embedded apps.
- Signs of Web 3.0: Rich, persistent streaming content – ubiquitous information

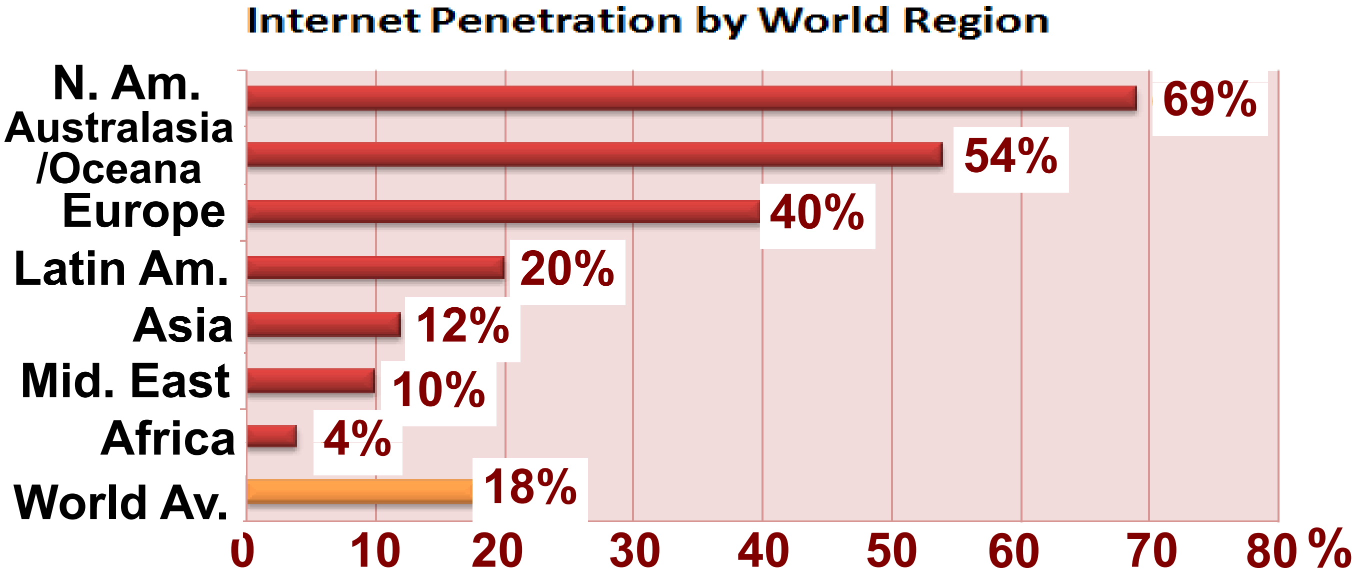

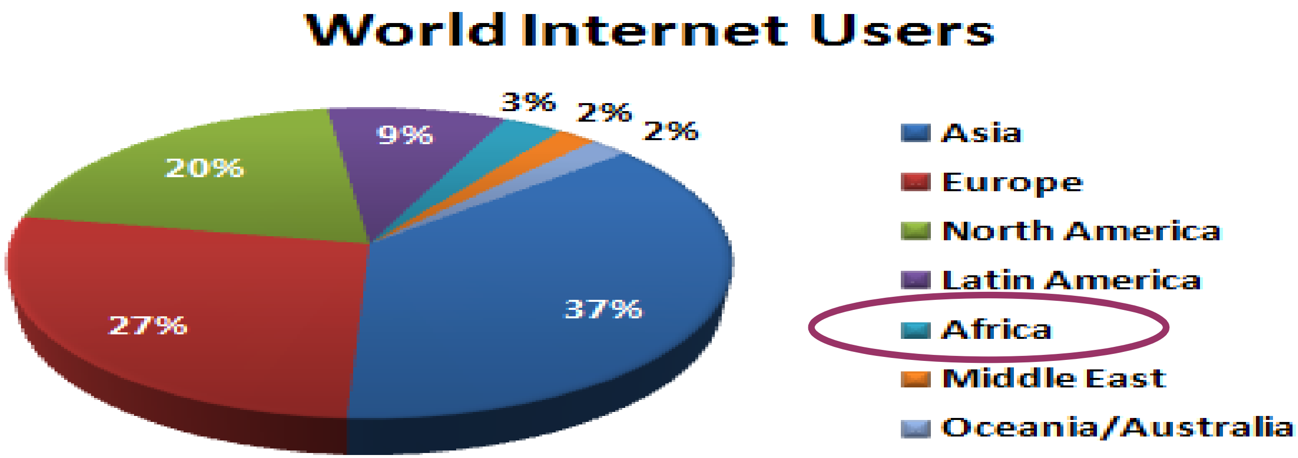

Broadband: 60M Each in the US and China

2^{nd} Revolution Network Drivers: Data Intensive Science Applications

- **High Energy & Nuclear Physics, AstroPhysics Sky Surveys: TByte to PByte "block" transfers at 1-10+ Gbps**
- **eVLBI: Many real time data streams at 1-10 Gbps**
- **BioInformatics, Clinical Imaging: GByte images on demand**
- **Fusion Energy: Time critical burst-data distribution; Distributed plasma simulations, visualization, analysis**
- ***Analysis Challenge:* Harness global computing, storage and *Network* resources, to *enable a global community to work collaboratively over great distances***

The LHC Data "Grid Hierarchy" Evolved: MONARC ➡ CMS & ATLAS Models, DISUN

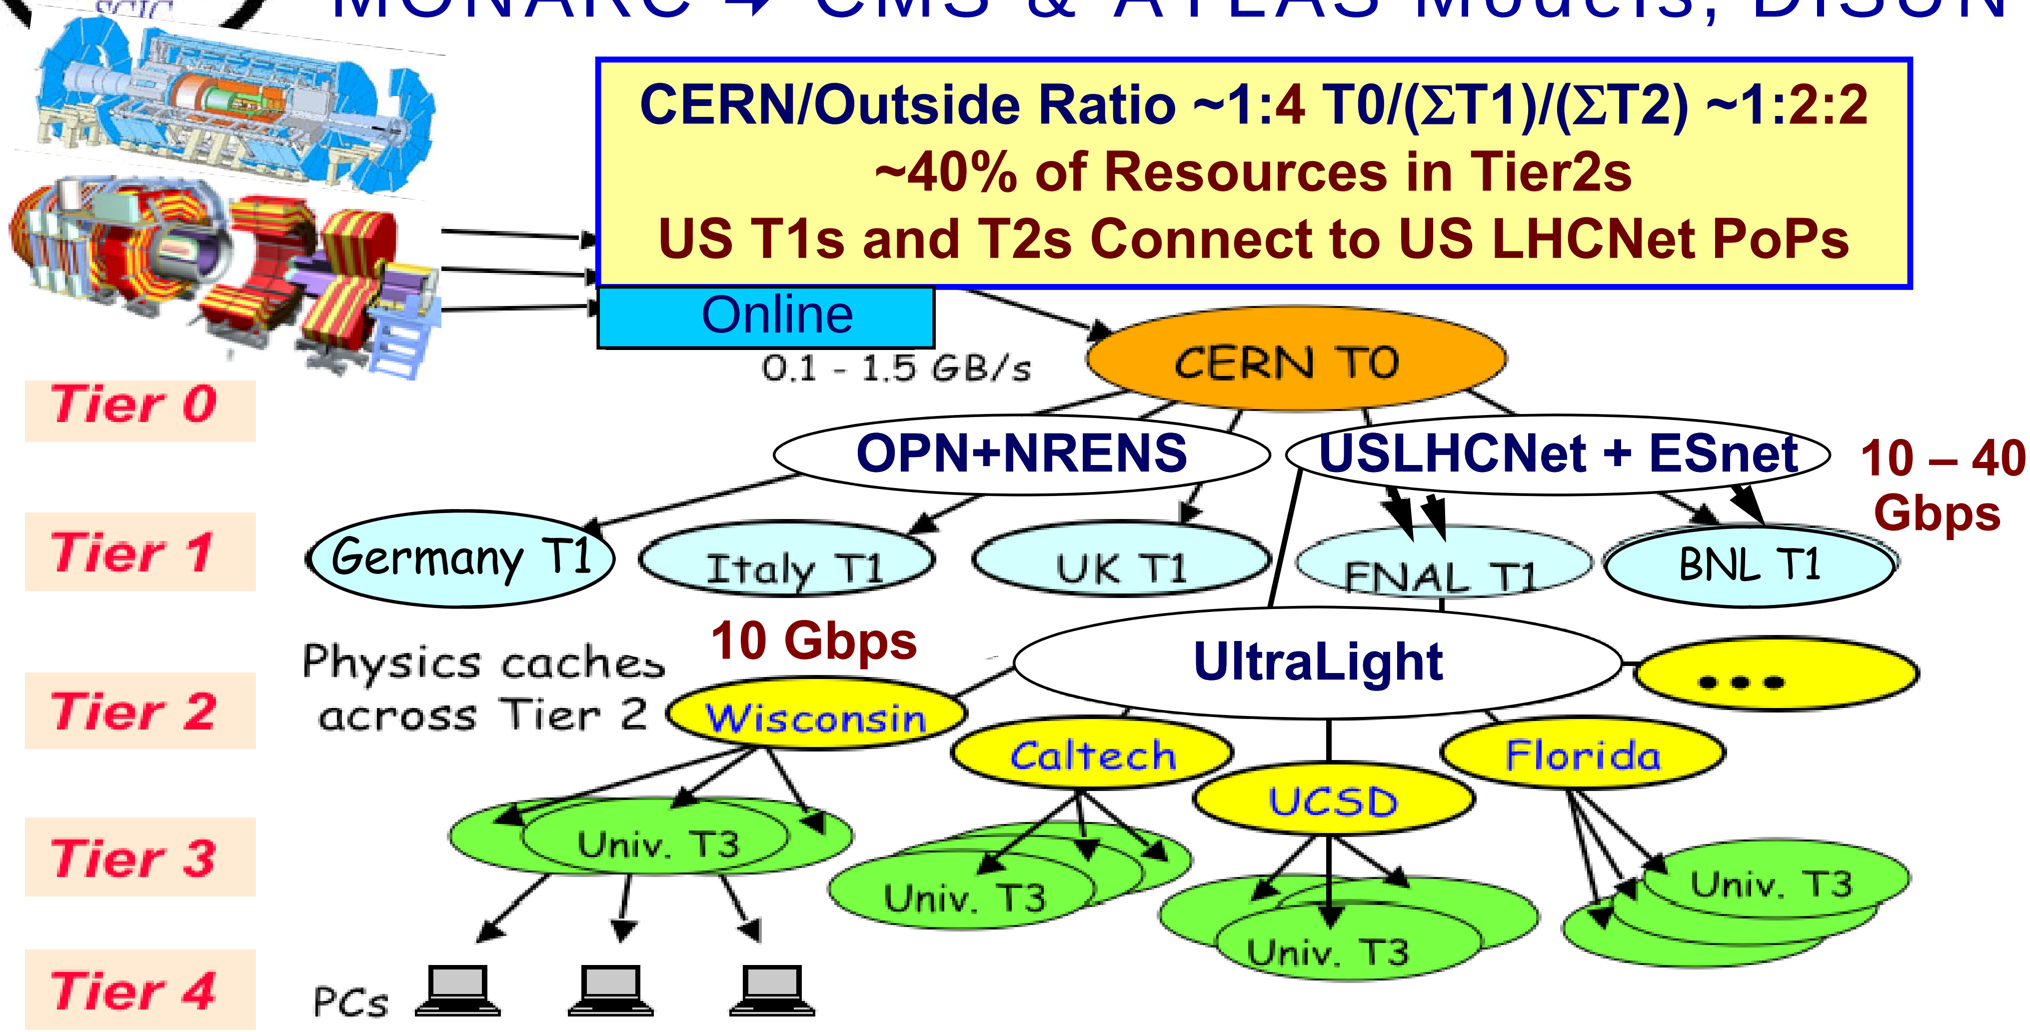

Outside/CERN Ratio Larger; Expanded Role of Tier1s & *Tier2s: Greater Reliance on Networks*

US CMS Tier1-Tier2 Network Utilization (January 2007)

- Develop and scale up production data movement services
- Acquire an understanding of how entire system performs under load

- **US CMS Tier1 (FNAL) to Tier2 (7 Sites) Traffic Disk-to-Disk**
- **Reached 9.5+ Gbps FNAL-Starlight; Saturated 10G Link**
- **2nd 10Gbps Link to Starlight deployed, with US LHCNet & ESnet**

`Daily' Graph (5 Minute Average)

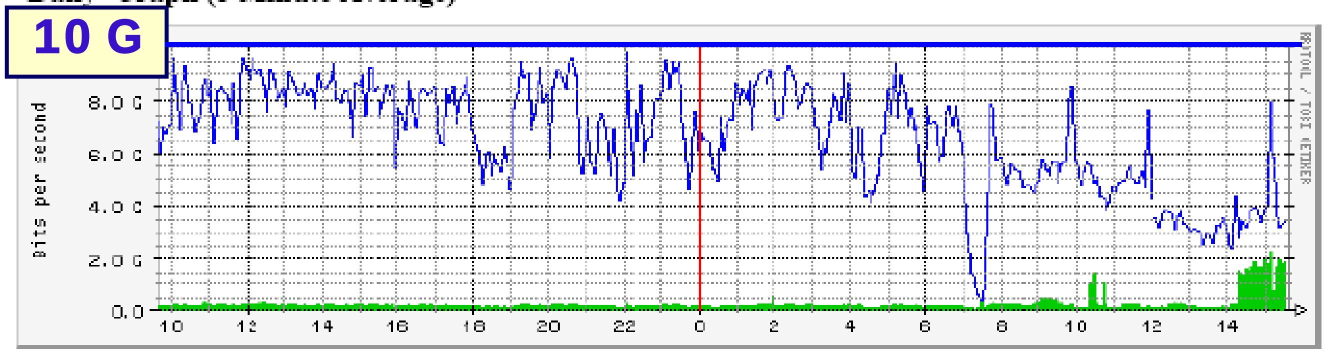

Max In: 2240.7 Mb/s (11.2%) Average In: 230.9 Mb/s (1.2%) Current In: 1894.4 Mb/s (9.5%)
Max Out: 9791.6 Mb/s (49.0%) Average Out: 6632.9 Mb/s (33.2%) Current Out: 3420.3 Mb/s (17.1%)

`Weekly' Graph (30 Minute Average)

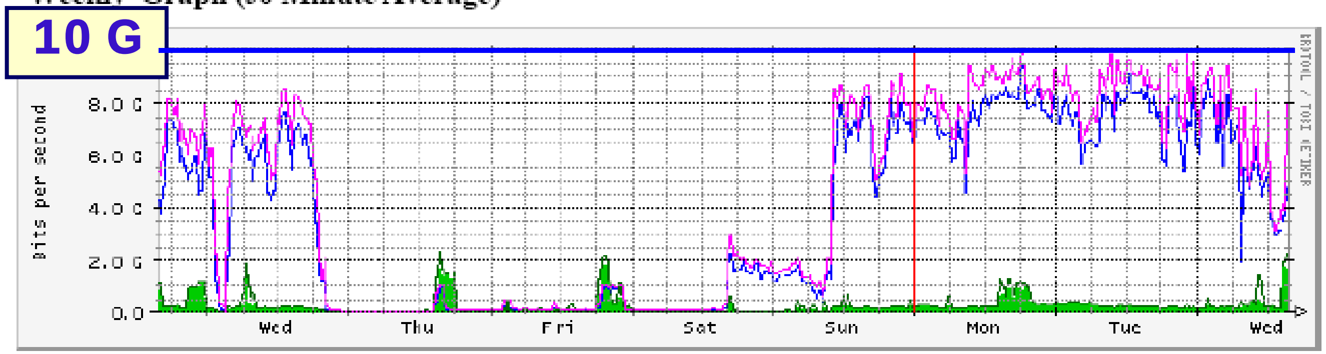

Max In: 1824.4 Mb/s (9.1%) Average In: 188.3 Mb/s (0.9%) Current In: 1691.6 Mb/s (8.5%)
Max Out: 9478.0 Mb/s (47.4%) Average Out: 3788.2 Mb/s (18.9%) Current Out: 4732.1 Mb/s (23.7%)

Accumulated data (Terabytes) received by CMS Data Centers ("tier1" sites) and many analysis centers ("tier2" sites) during the past four months (8 petabytes of data) [LHC/CMS]

This sets the scale of the LHC distributed data analysis problem.

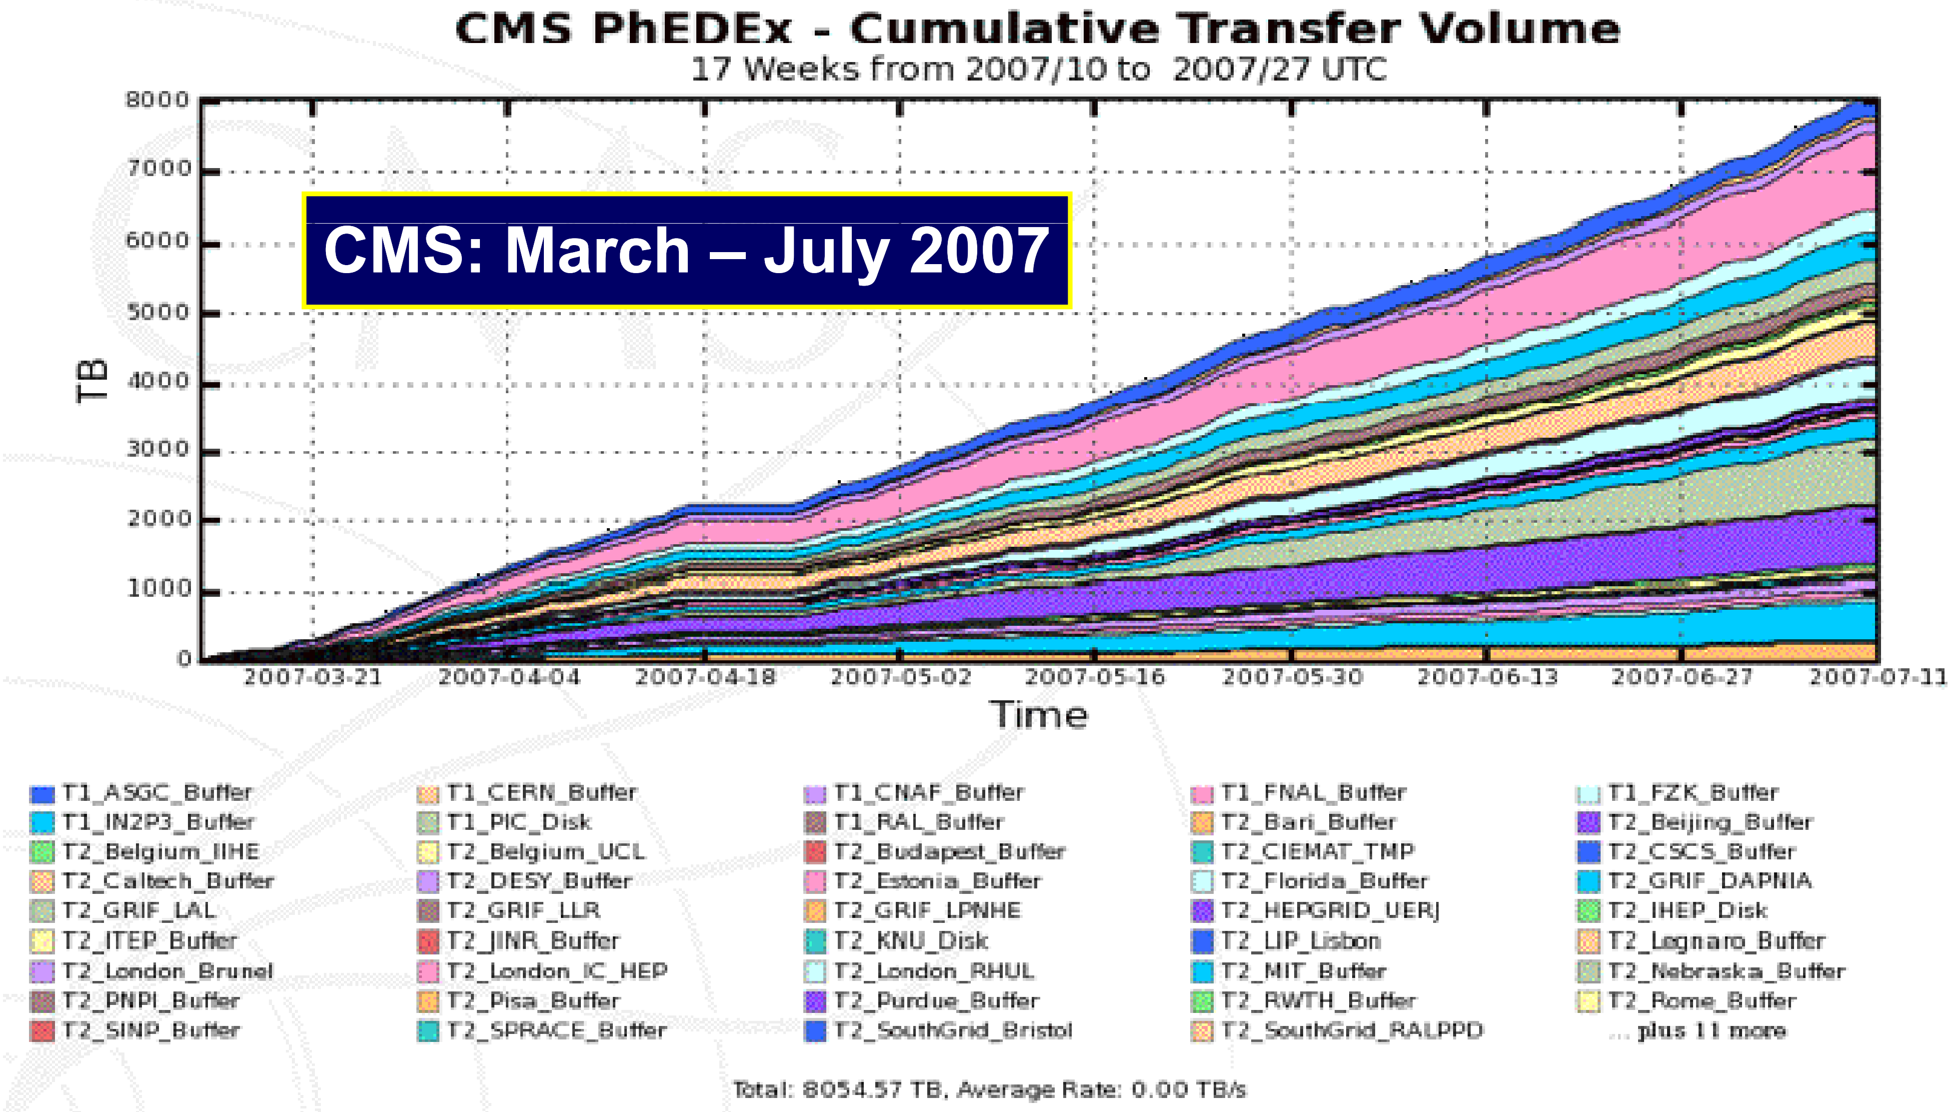

PhEDEx: CMS Data Transfers (8/29/07)

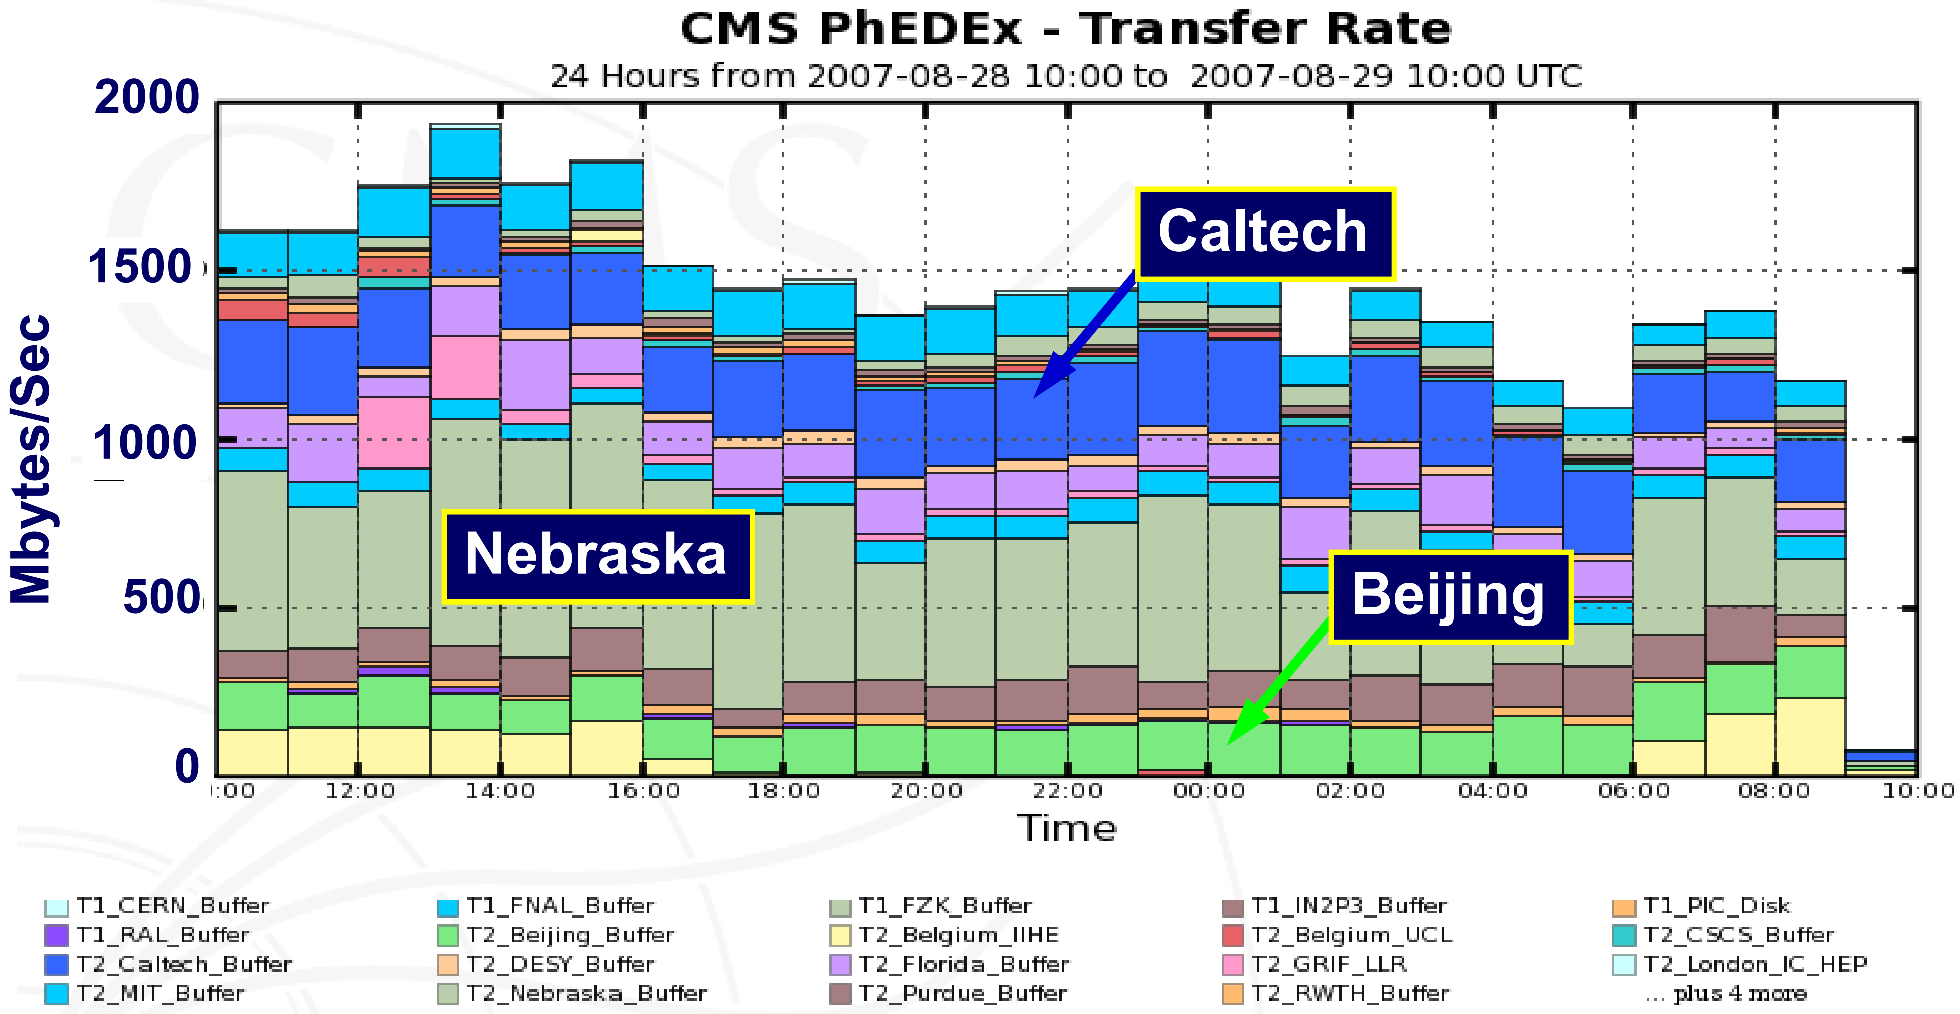

Total to 1.9 Gbytes/sec. Nebraska Tier2 to 700 Mbytes/sec

Large-Scale Science is Beginning to Dominate all Traffic

W. Johnston ESnet

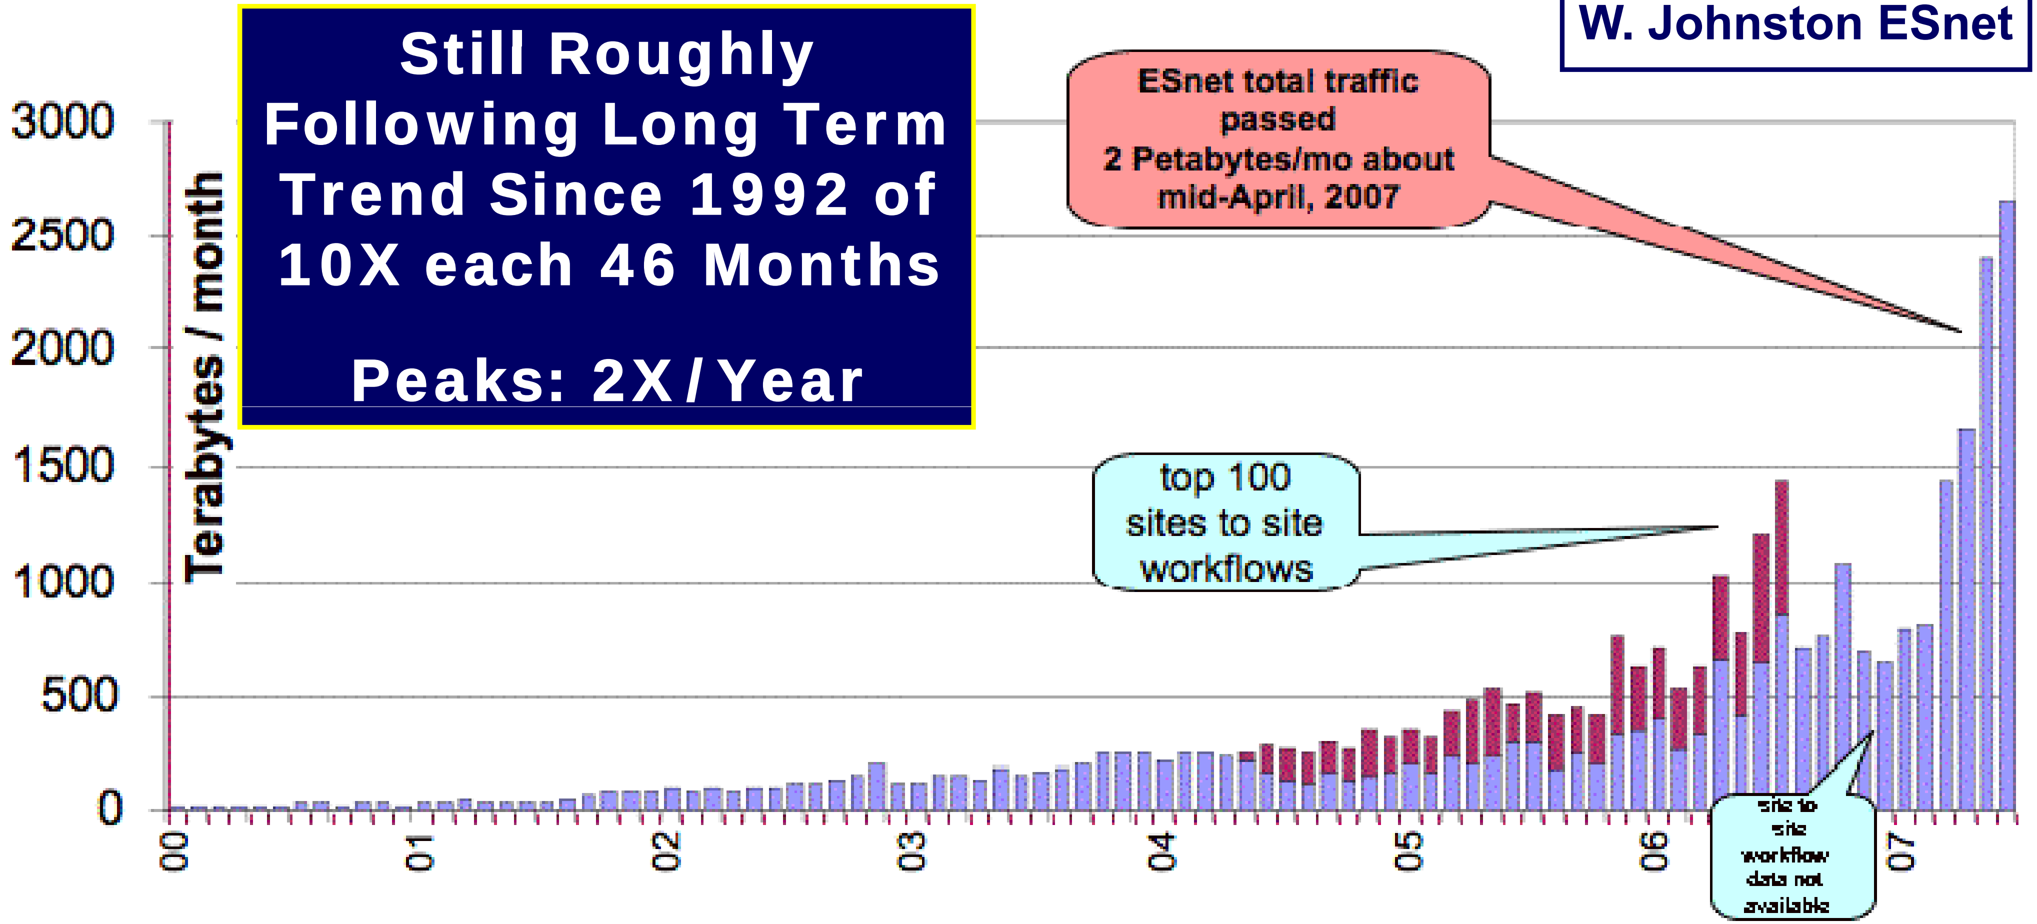

ESnet Monthly Accepted Traffic, January, 2000 – May, 2007

- ESnet is currently transporting more than1 petabyte (1000 terabytes) per month
- More than 50% of the traffic is now generated by the top 100 sites ⇒ large-scale science dominates all ESnet traffic

LHC "Optical Private Network" (OPN) (Status August 2007)

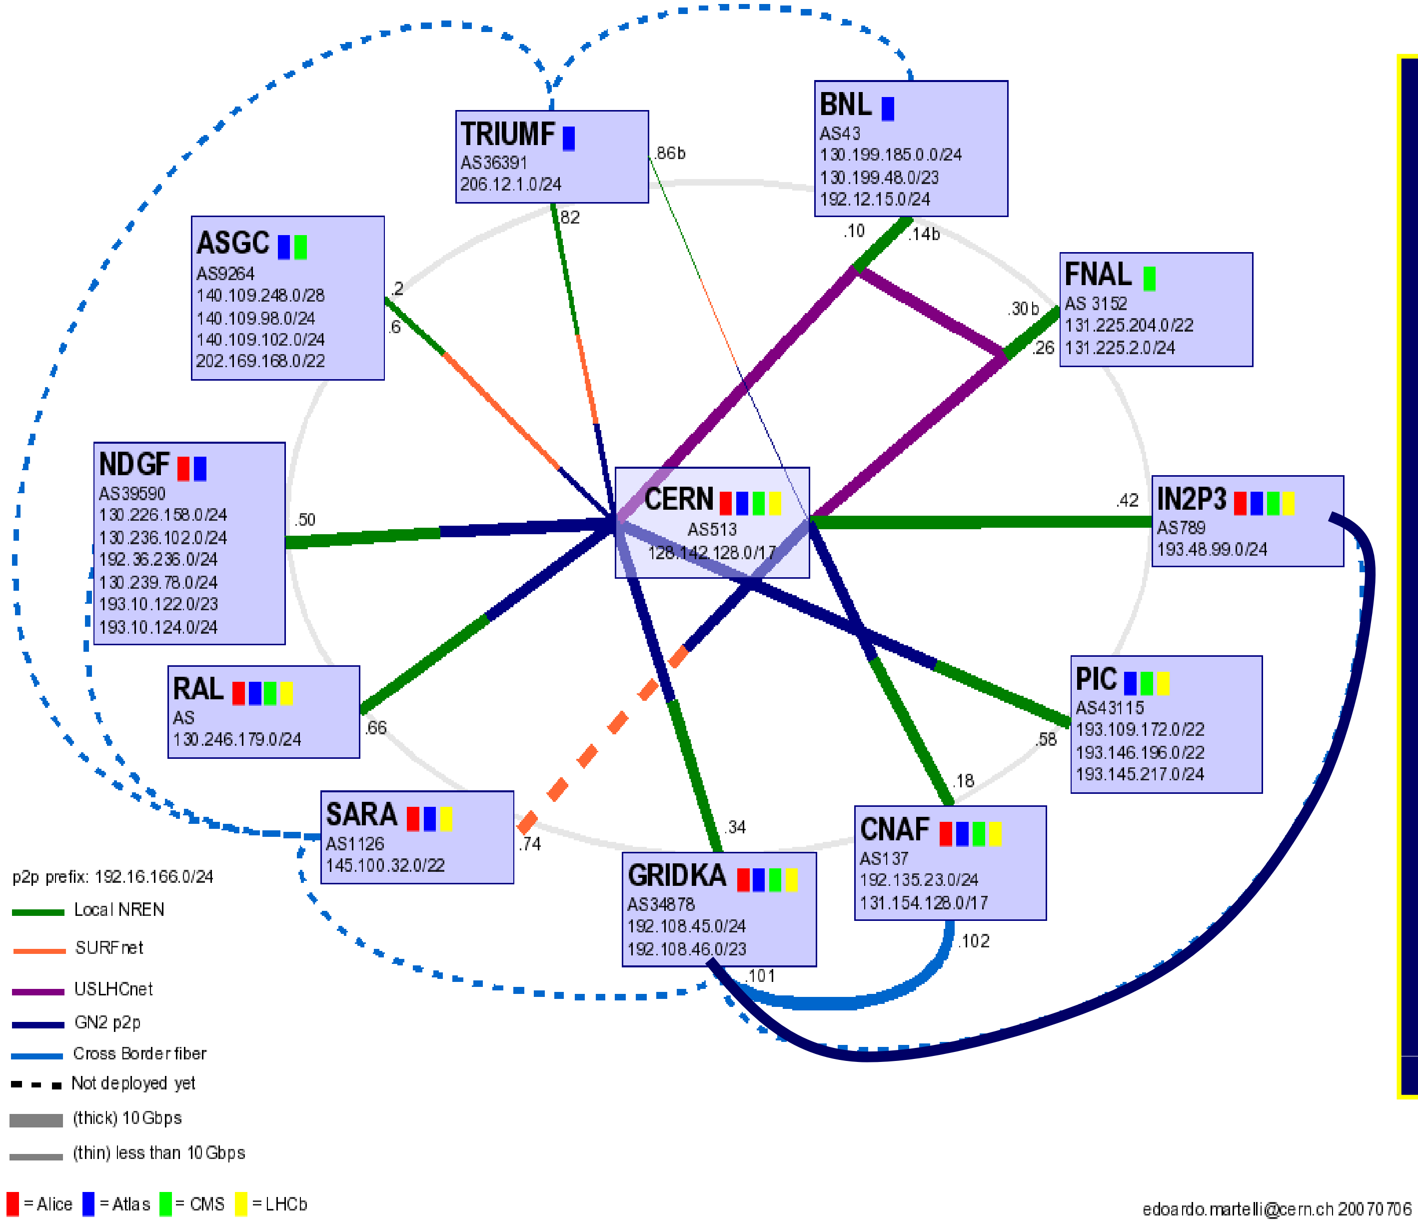

Some Tier1-Tier1 connectivity not deployed yet; needed for resilience.

Progressing: e.g. 10G GridKa-IN2P3 link Sept. 1

US LHCNet & ESnet Plan 2007-2010: 20-80Gbps US-CERN, ESnet MANs
Funding Issue

US-LHCNet: Wavelength Quadrangle NY-CHI-GVA-AMS
2007-10: 30, 40, 60, 80 Gbps

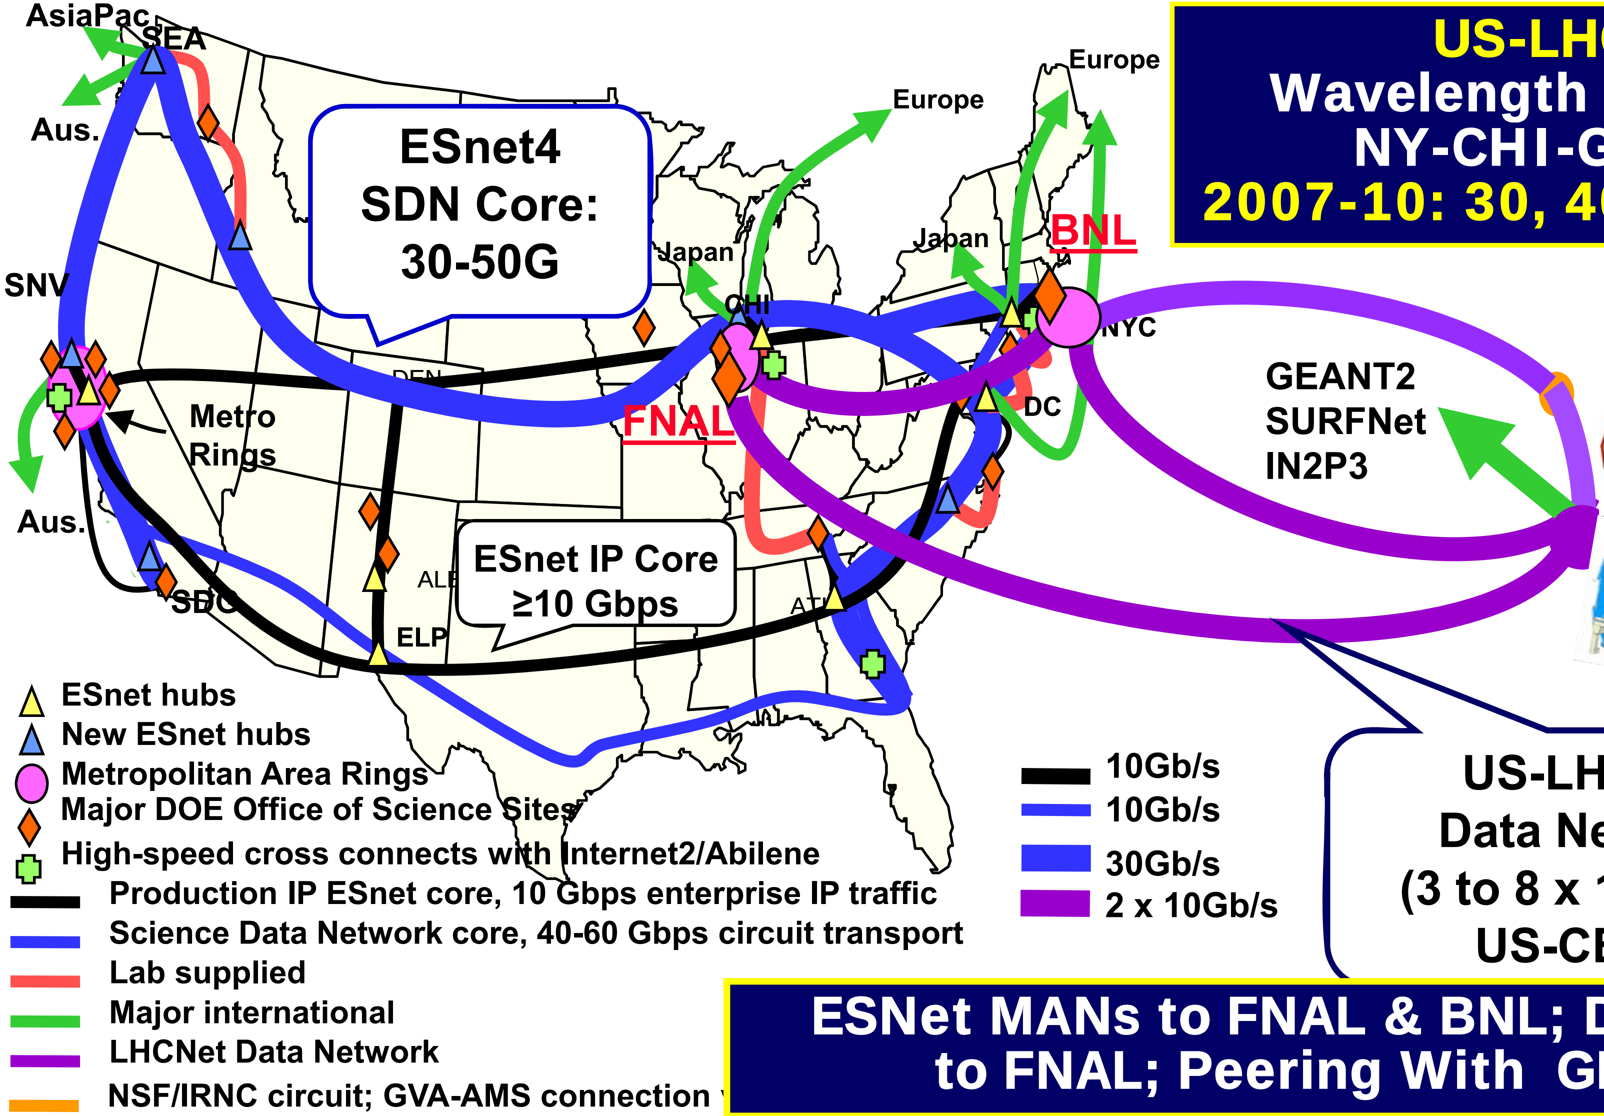

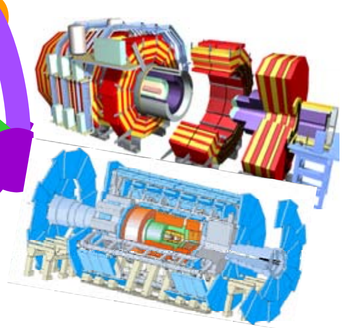

US-LHCNet Data Network (3 to 8 x 10 Gbps US-CERN)

ESNet MANs to FNAL & BNL; Dark Fiber to FNAL; Peering With GEANT

Updated HEP Bandwidth Roadmap for Major Links (in Gbps)

Year	Production	Experimental	Remarks
2001	0.155	0.622-2.5	SONET/SDH
2002	0.622	2.5	SONET/SDH DWDM; GigE Integ.
2003	2.5	10-20	DWDM; 1 + 10 GigE Integration
2005	10-20	2-10 X 10	λ Switch; λ Provisioning
2007	3-4 X 10	~10 X 10; 100 Gbps	1st Gen. λ Grids
2009	~6 X 10 or 100	~20 X 10 or ~2 X 100	100 Gbps λ Switching
2011	~20 X 10 or 2 X 100	~10 X 100	2nd Gen λ Grids Terabit Networks
2013	~Terabit	~MultiTbps	~Fill One Fiber

Continuing Trend: ~400-1000 Times Bandwidth Growth Per Decade
Paralleled by ESnet Roadmap for Data Intensive Sciences

ESnet4

50-60 Gbps by 2009-10; 500-600 Gbps ~2010-12

W. Johnston
ESnet

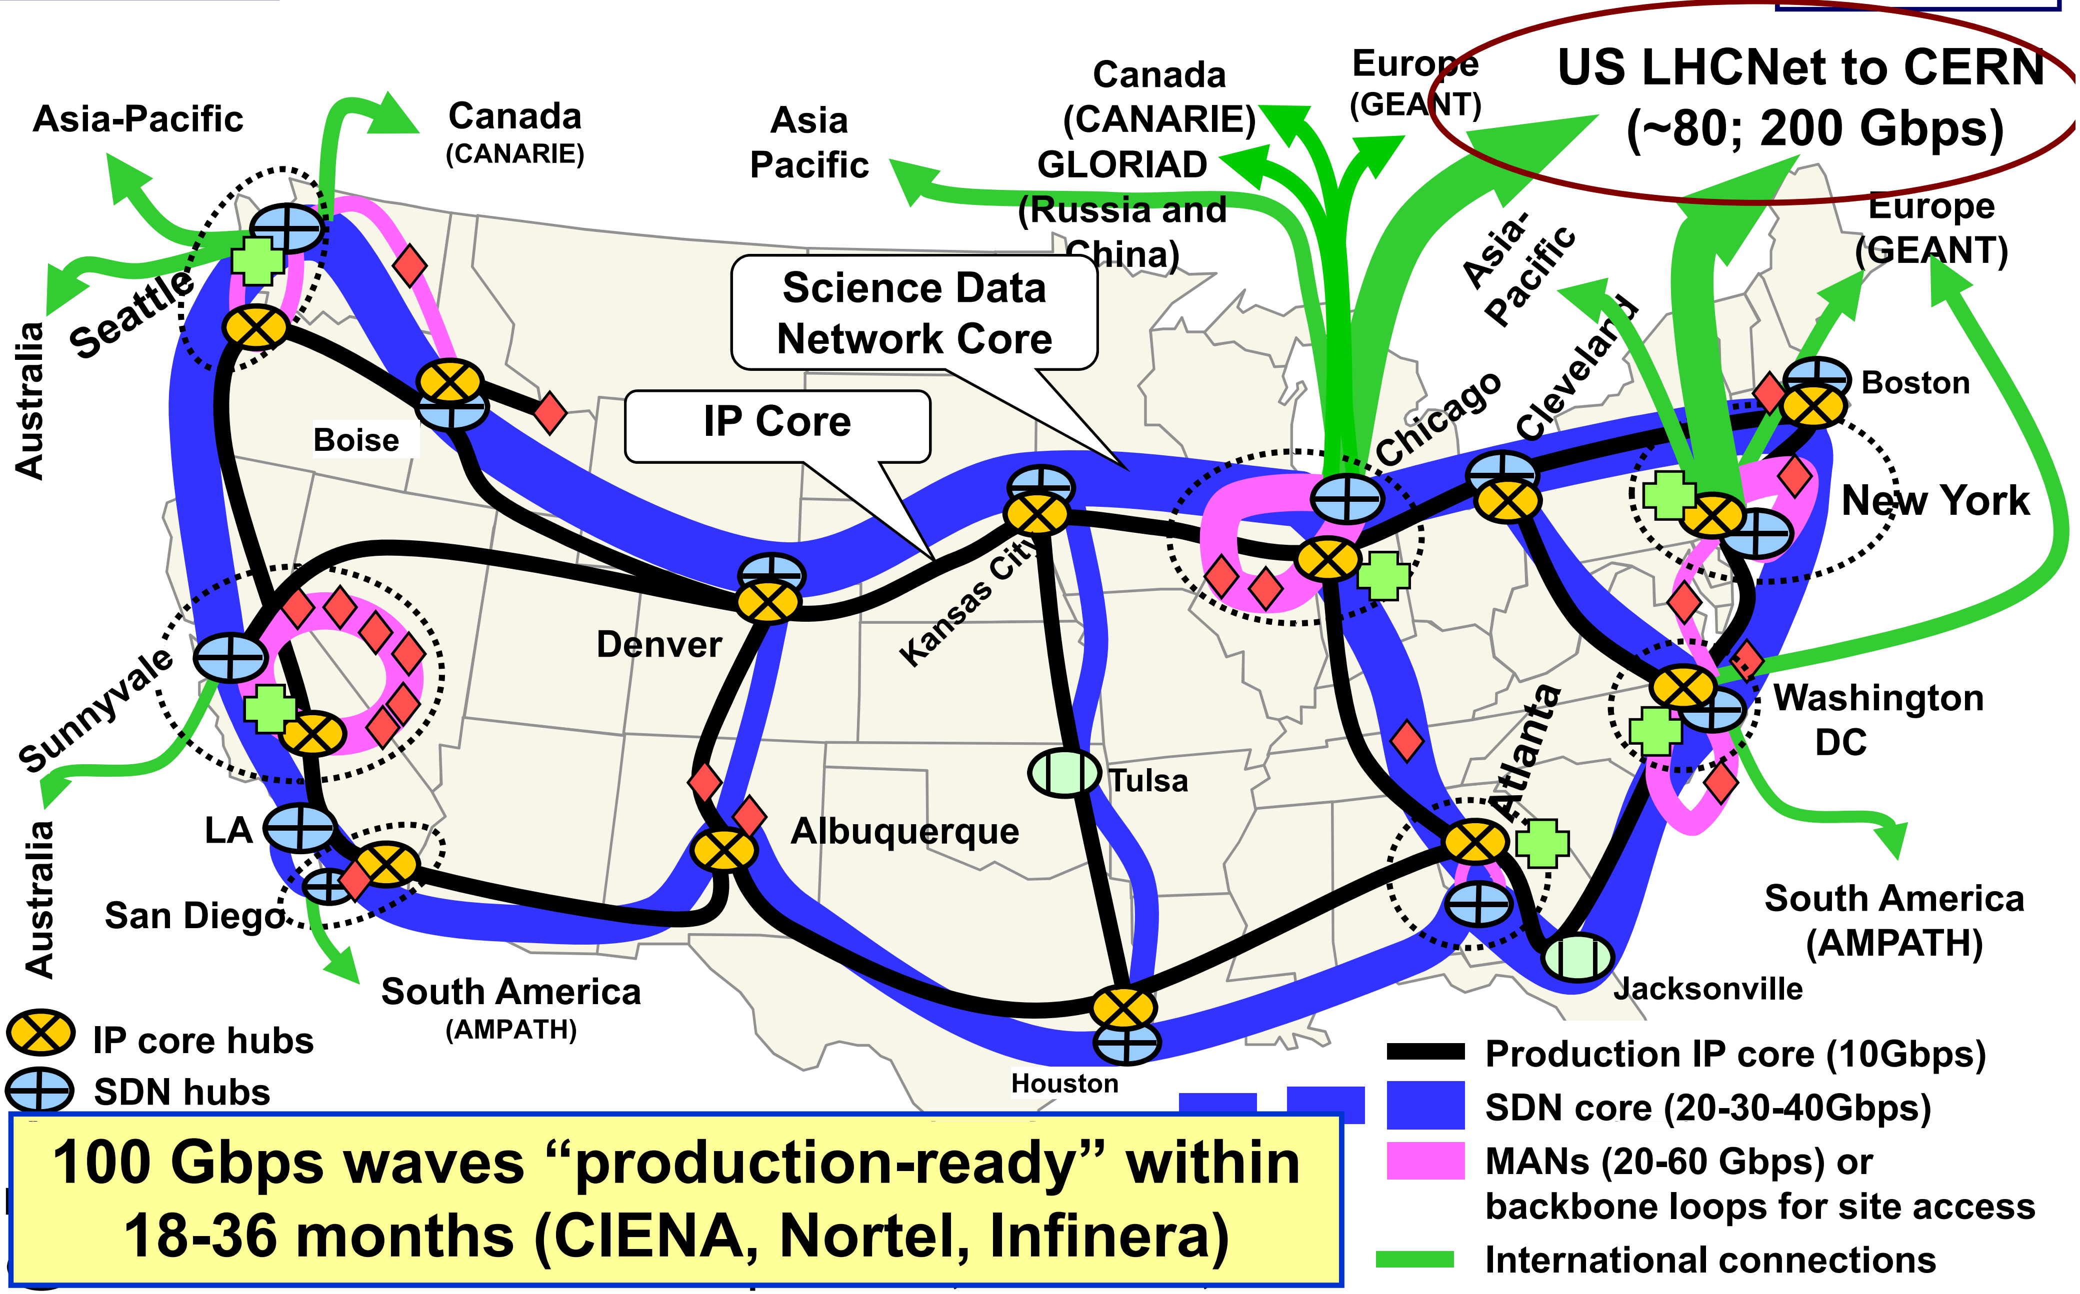

100 Gbps waves "production-ready" within 18-36 months (CIENA, Nortel, Infinera)

2nd Revolution: Networks for Research & Education and Data Intensive Science

- Current generation of 10 Gbps R&E backbones arrived in 2001-5 in US, Europe, Japan, Korea; *China. Transition to N X 10G*
  - ➔ Bandwidth Growth: from 4 to 2500 Times in 5 Years; >> Moore's Law
- Proliferation of 10G links across the Atlantic & Pacific; Use of multiple 10G Links (e.g. US-CERN) along major paths began in Fall 2005
  - ➔ On track for >10 X 10G networking for LHC, in production by 2007-8
- Rapid Spread of "Dark Fiber" and DWDM: the emergence of Continental, Nat'l, State & Metro "Hybrid" Networks in Many Nations
  - ➔ Cost-effective N X 10G Backbones
  - ➔ Complemented by Point-to-point "Light-paths" for "Data Intensive Science", notably HEP
  - ➔ Outlook: 100G waves within the next 5 years; dynamic paths
- Technology evolution continues to drive performance higher, equipment costs Lower; HEP is a leading user/developer
  - ➔ Commodity GbE and 10 GbE NICs on servers; multi- core CPUs; PCI Express bus for high speed data transport

➔ 2007 Outlook: *Continued growth in bandwidth deployment & use*

GÉANT2: Consortium of 34 NRENs

22 PoPs, ~200 Sites
38k km Leased Services, 12k km Dark Fiber
Supporting Light Paths for *LHC*, *eVLBI*, et al.

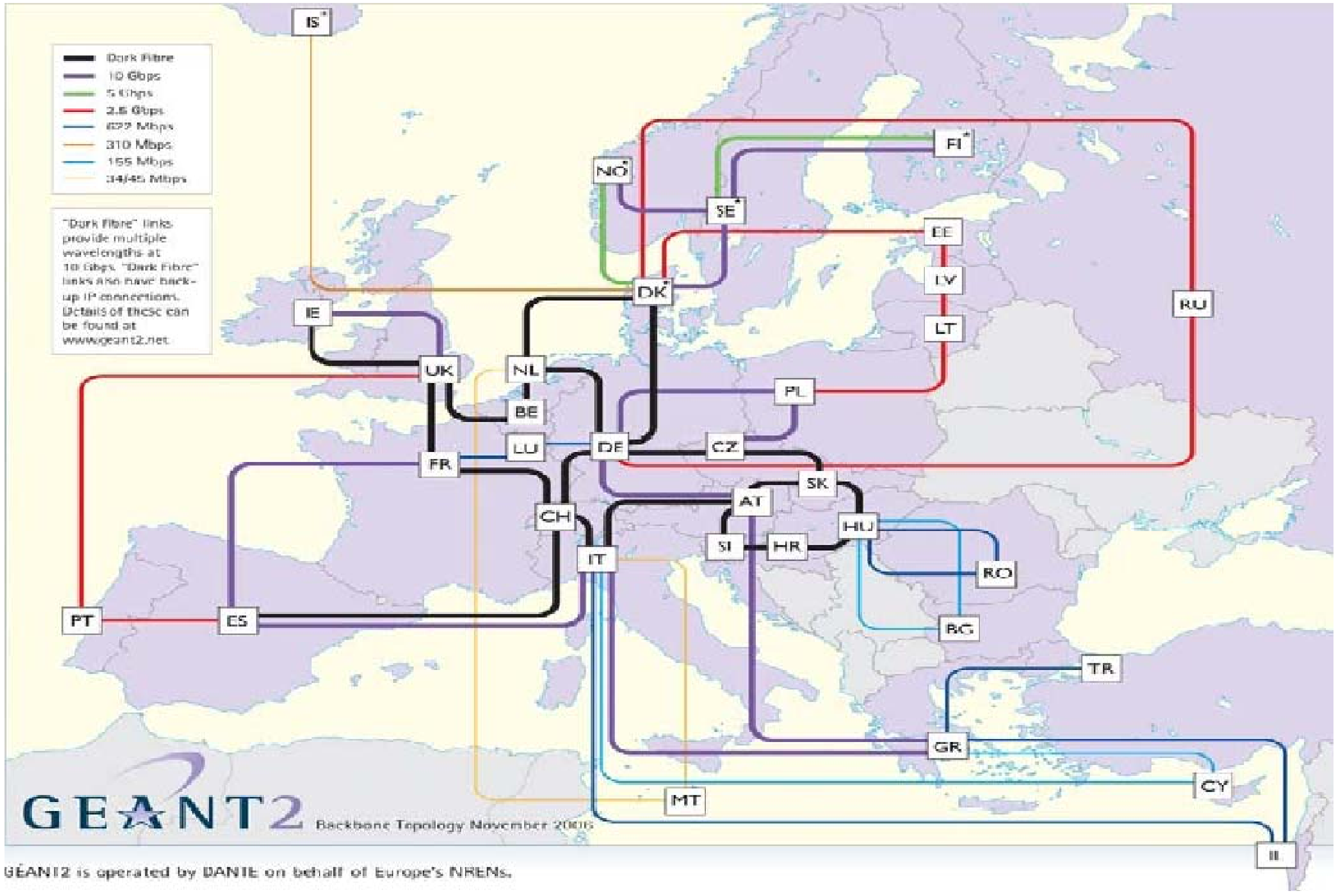

Dark Fiber Core Among <u>16 Countries:</u>

- Austria
- Belgium
- Bosnia-Herzegovina
- Czech Republic
- Denmark
- France
- Germany
- Hungary
- Ireland
- Italy,
- Netherland
- Slovakia
- Slovenia
- Spain
- Switzerland
- United Kingdom

Multi-Wavelength Core (to 40λ) + 0.6-10G Loops

H. Doebbeling

Pioneering "Light Paths": Canada (CANARIE) & Netherlands (SURFNet)

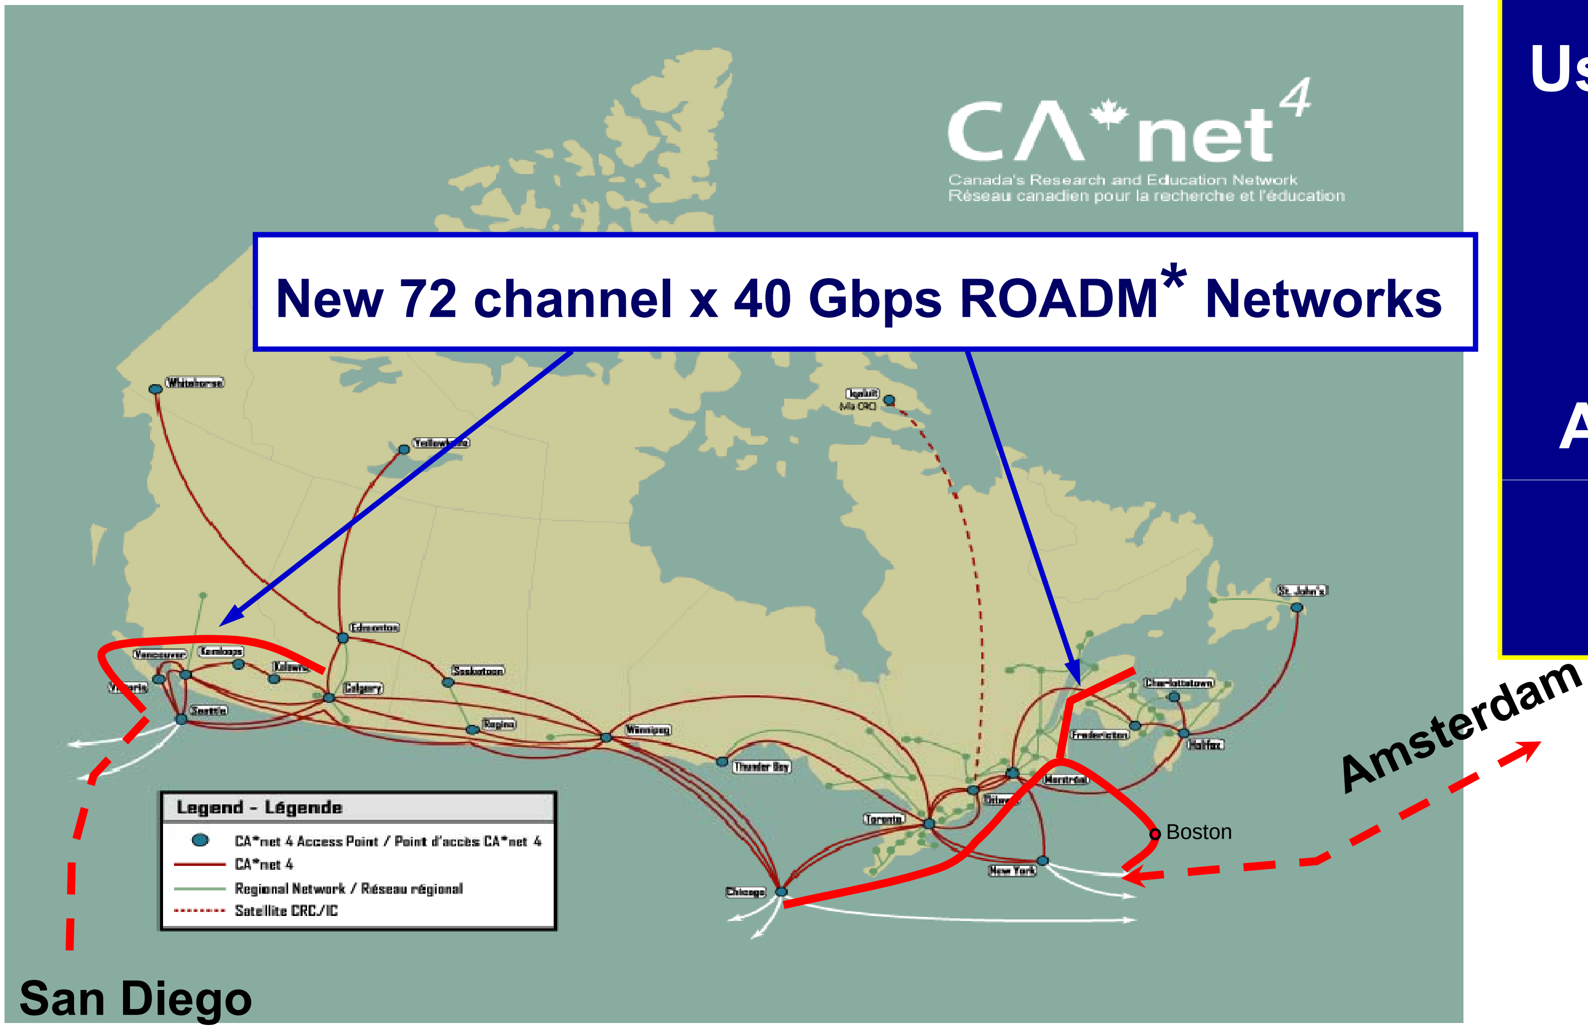

User-Controlled Light Paths

72λ Core: 40G Waves

Also connects to U.S. and Europe

W. St. Arnaud

*Reconfigurable Optical Add-Drop Multiplexer

GLIF Open Lambda Exchanges (GOLE)

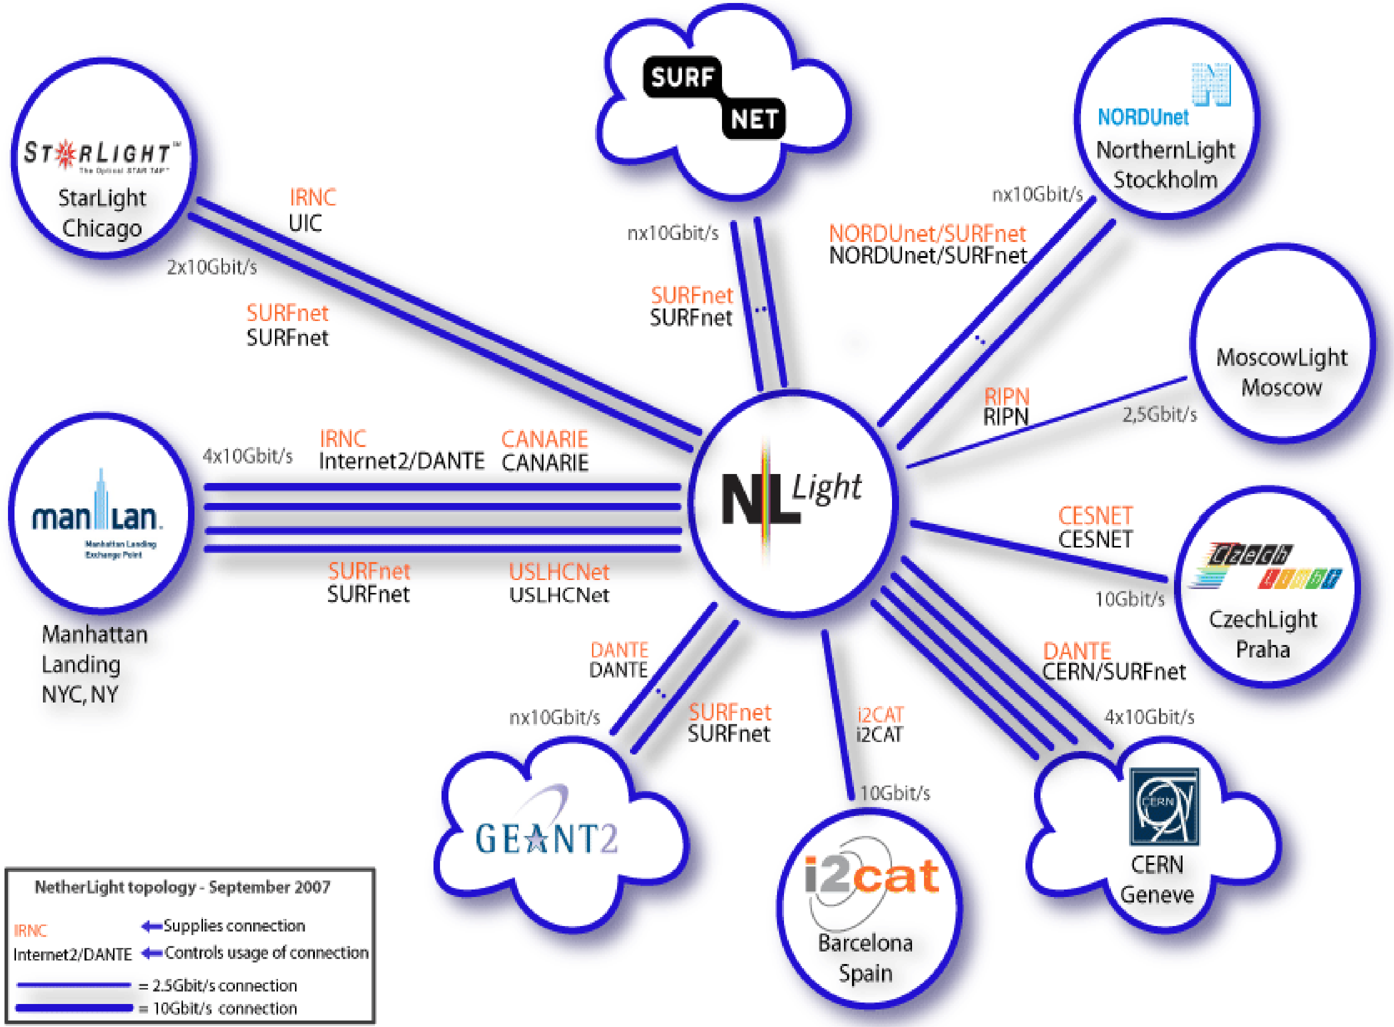

- AMPATH - Miami
- CERN/Caltech – Geneva+U.S.
- CzechLight - Prague
- HKOEP - Hong Kong
- KRLight - Daejoen
- MAN LAN - New York
- MoscowLight - Moscow
- NetherLight - Amsterdam
- NGIX-East – Wash. D.C.
- NorthernLight - Stockholm
- Pacific Wave (L.A.)
- Pacific Wave (Seattle) - Pacific Wave (Sunnyvale)
- StarLight - Chicago
- T-LEX - Tokyo
- UKLight - London

Internet2's New Backbone with "Light Paths"

Internet2-NLR Merger Discussions

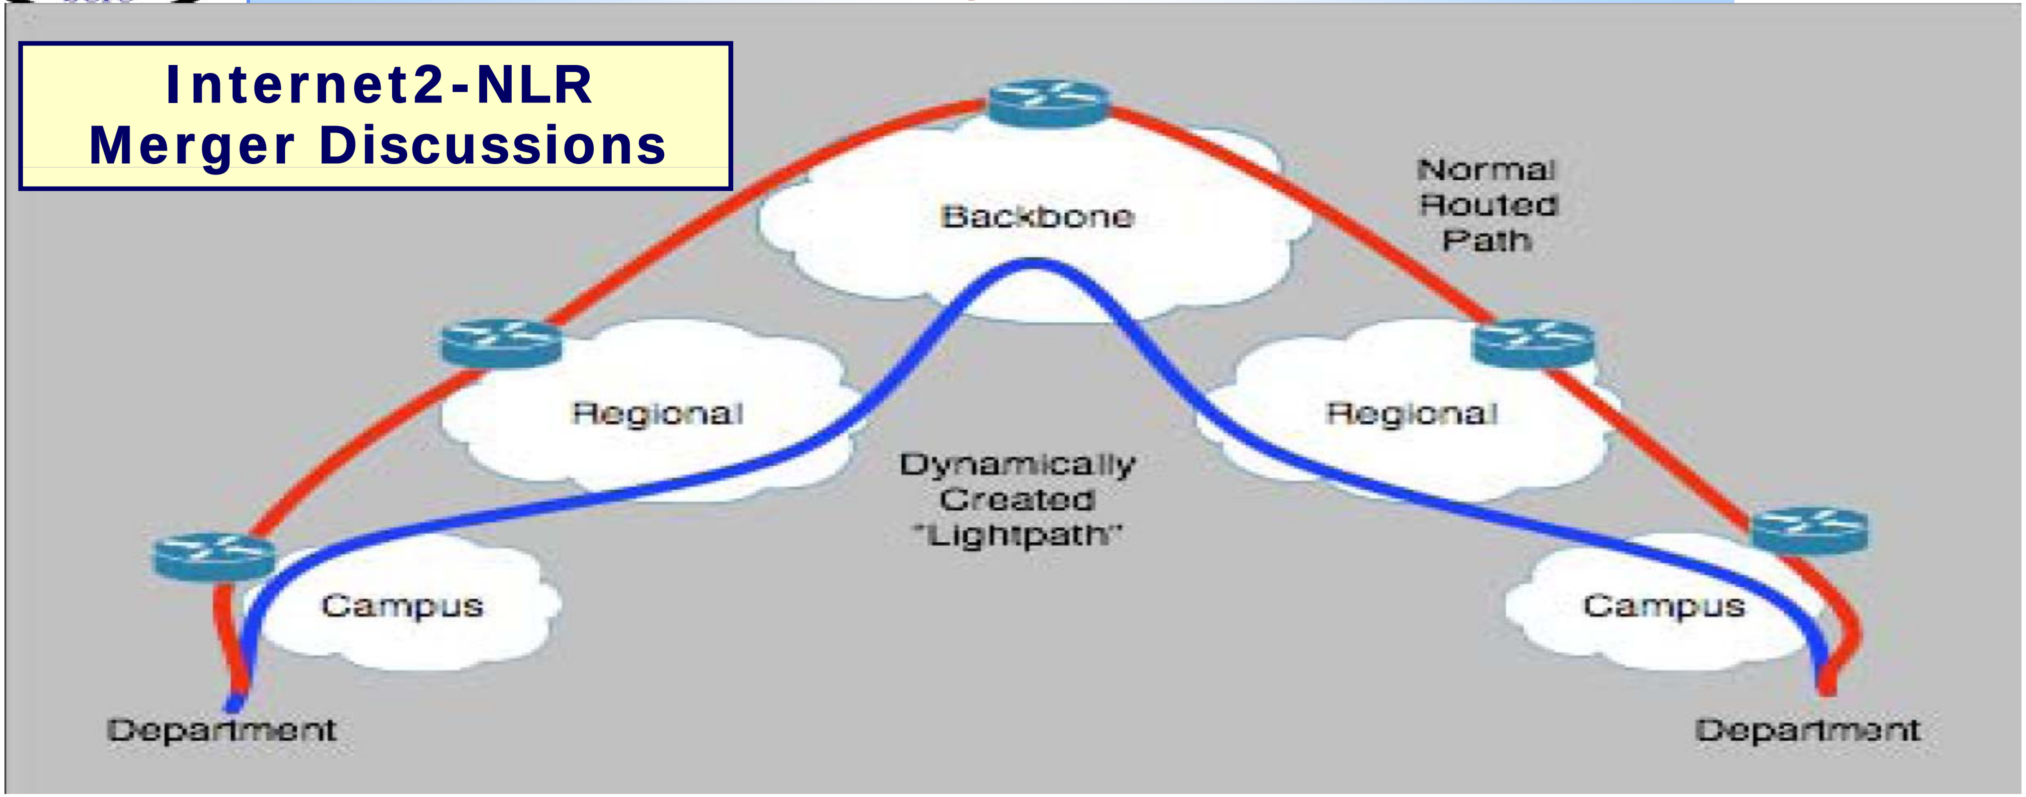

- Initial deployment – 10 x 10 Gbps wavelengths over the footprint
- First round maximum capacity – 80 x 10 Gbps wavelengths; expandable
- Scalability – potential migration to 40 Gbps or 100 Gbps capability
- Transition underway (since 10/2006), until end of 2007

+Paralleled by Initiatives in: nl, ca, jp, uk, kr; *pl, cz, sk, pt, ei, gr, hu, si, lu, no, is, dk* ... + >30 US states

ICFA SCIC

GLORIAD: An Optical Ring Around Earth

http://www.gloriad.org

GLORIAD in 2005-7

- 2.5G Hong Kong – Seattle
- 10G Daejon (Korea) – Seattle
- 10G Stockholm – Moscow
- 10G Amsterdam – Moscow
- 3 X 10G: SURFnet & IRNC, Amsterdam - US

1st 1 Gigabit Lightpath

- Between Moscow (KIAE) & Chicago (NCDM)
- 1.7 TBytes in 4.5 Hrs

2nd Quiet Revolution in Science and R&E Networks Continues

- ***2000-2007: HEP, working with computer scientists and network engineers, has developed the knowledge to use long distance networks efficiently, at high occupancy, for the first time***
  - **"Demystification" of large long range data flows with TCP:**
  - **Exploit advances in the TCP stack (e.g. FAST TCP), Linux Kernel (2.6.20), end system architecture, network interfaces and drivers**
- **Just one to a few server-pairs with 10 GbE interfaces match a 10 Gbps Link**
- **Make the advances *widely* accessible**

Bandwidth Challenge at SC2005:

Caltech, CERN, FNAL, BNL, SLAC, UM, UF, ESNet, I2…

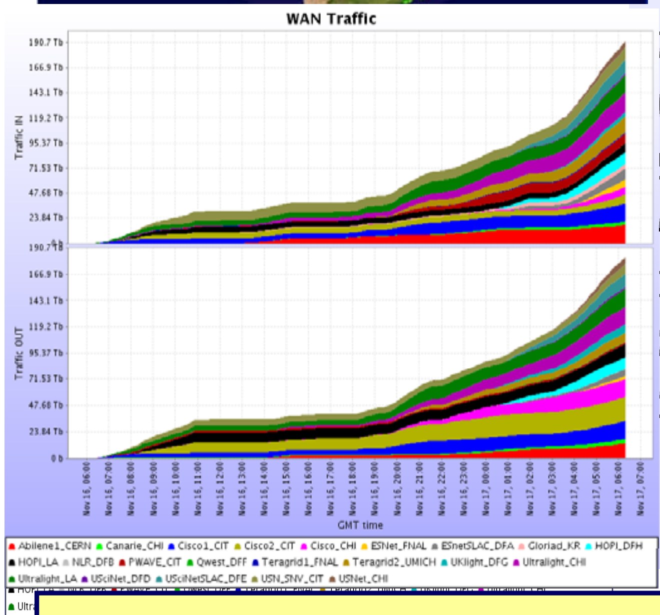

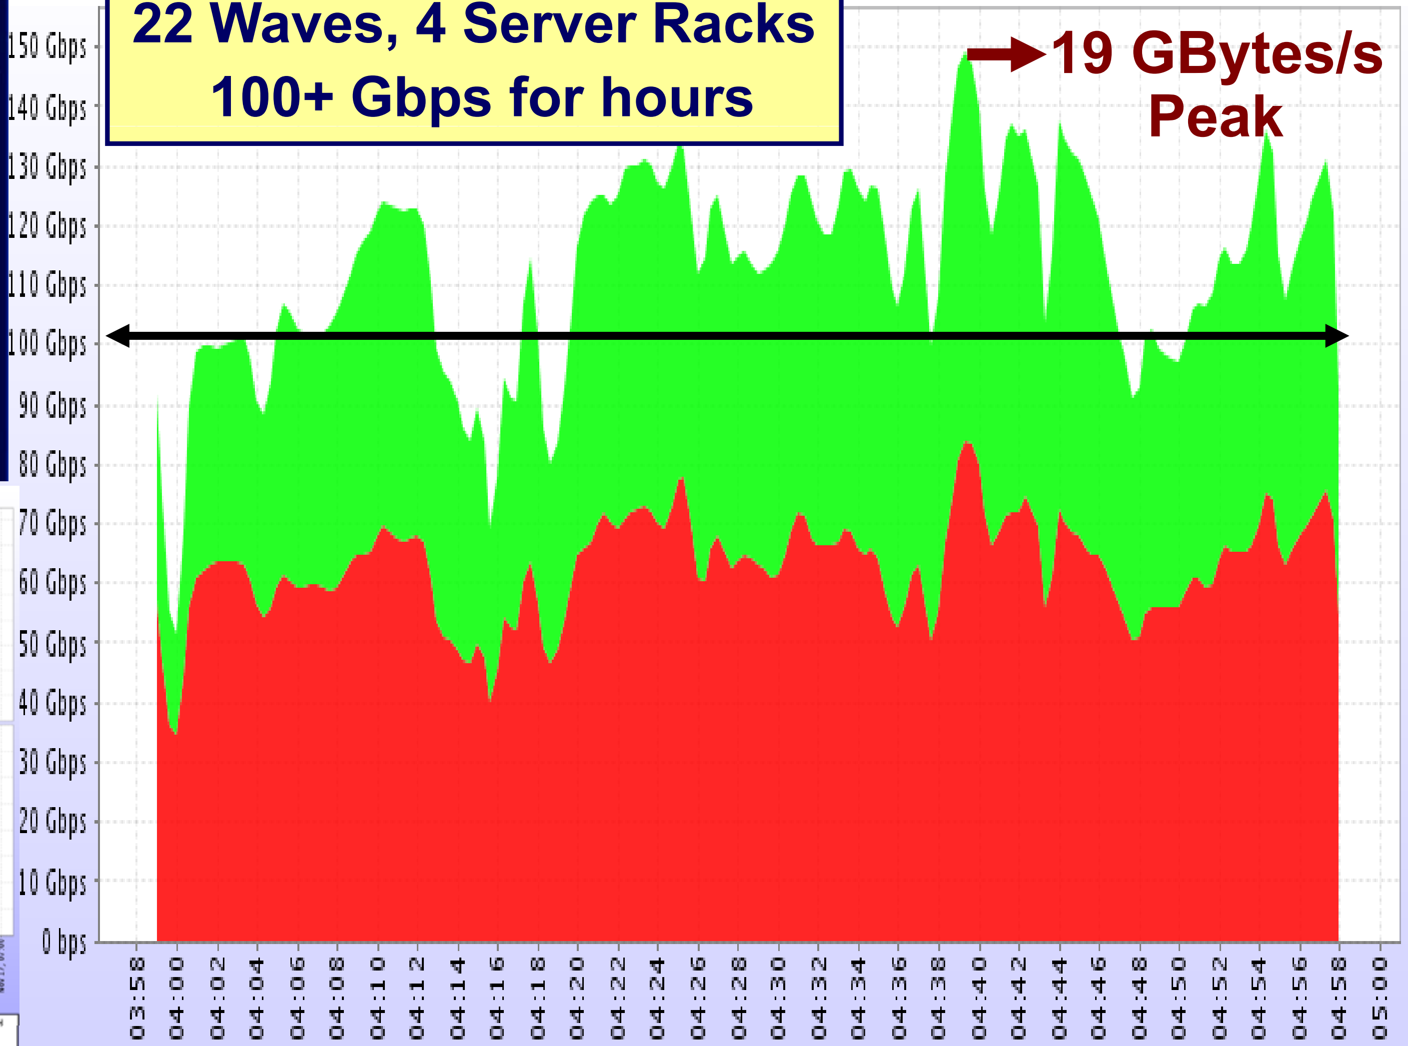

475 TB Total in < 24h; Sustained Rate of 1.1 Petabyte Per Day

FDT – Fast Data Transfer

I. Legrand

An easy-to-use application for efficient data transfers

- Written in Java (with NIO libraries):
  runs on all major platforms
- Uses asynch. multithreaded system to achieve smooth, linear data flow:
  - stream a dataset (list of files) continuously, through an open TCP socket
    - ➡ protocol does Not Stop between files
  - send buffers at a rate matched to the monitored capability of end to end path
  - use independent threads to read & write on each physical device
  - use appropriate size of buffers for disk I/O and networking
- Ckeck-pointing: Resumes a file transfer session, if interrupted
- Can "plug-in" external security APIs and use them to authenticate and authorize clients: SSH, GSI-SSH, Globus-GSI, SSL

- SC06 BWC: Stable disk-to-disk flows Tampa-Caltech: 10-to-10 and 8-to-8 1U Server-pairs for 9 + 7 = 16 Gbps; then Solid overnight. Using One 10G link

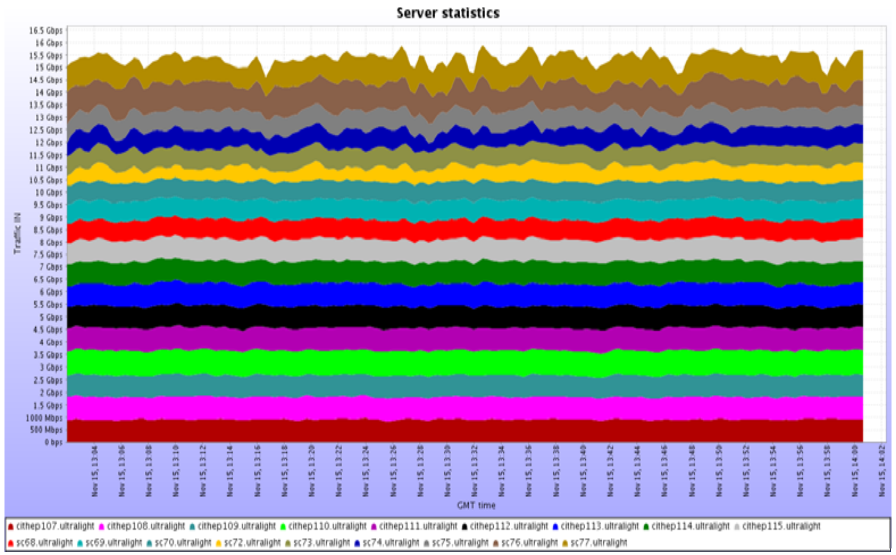

- 17.77 Gbps BWC peak;
  + 8.6 Gbps to and from Korea

New Capability Level: ~70-100 Gbps per rack of low cost 1U servers

What Networks Need to Do

W. Johnston (ESnet manager)

- The above examples currently only work in carefully controlled environments with the assistance of computing and networking experts
- For this essential approach to be successful in the long-term it must be routinely accessible to discipline scientists - without the continuous attention of computing and networking experts
- In order to
  - facilitate operation of multi-domain distributed systems
  - accommodate the projected growth in the use of the network
  - facilitate the changes in the types of traffic

  the architecture and services of the network must change

- **The general requirements for the new architecture are that it provide:**
  1) **Support the high bandwidth data flows of large-scale science including scalable, reliable, and very high-speed network connectivity to end sites**
  2) **Dynamically provision virtual circuits with guaranteed quality of service (e.g. for dedicated bandwidth and for traffic isolation)**
  3) **provide users and applications with meaningful monitoring end-to-end (across multiple domains)**

US LHCNet: Transition Now to CIENA Optical Muxes

Emerging Standards VCAT, LCAS

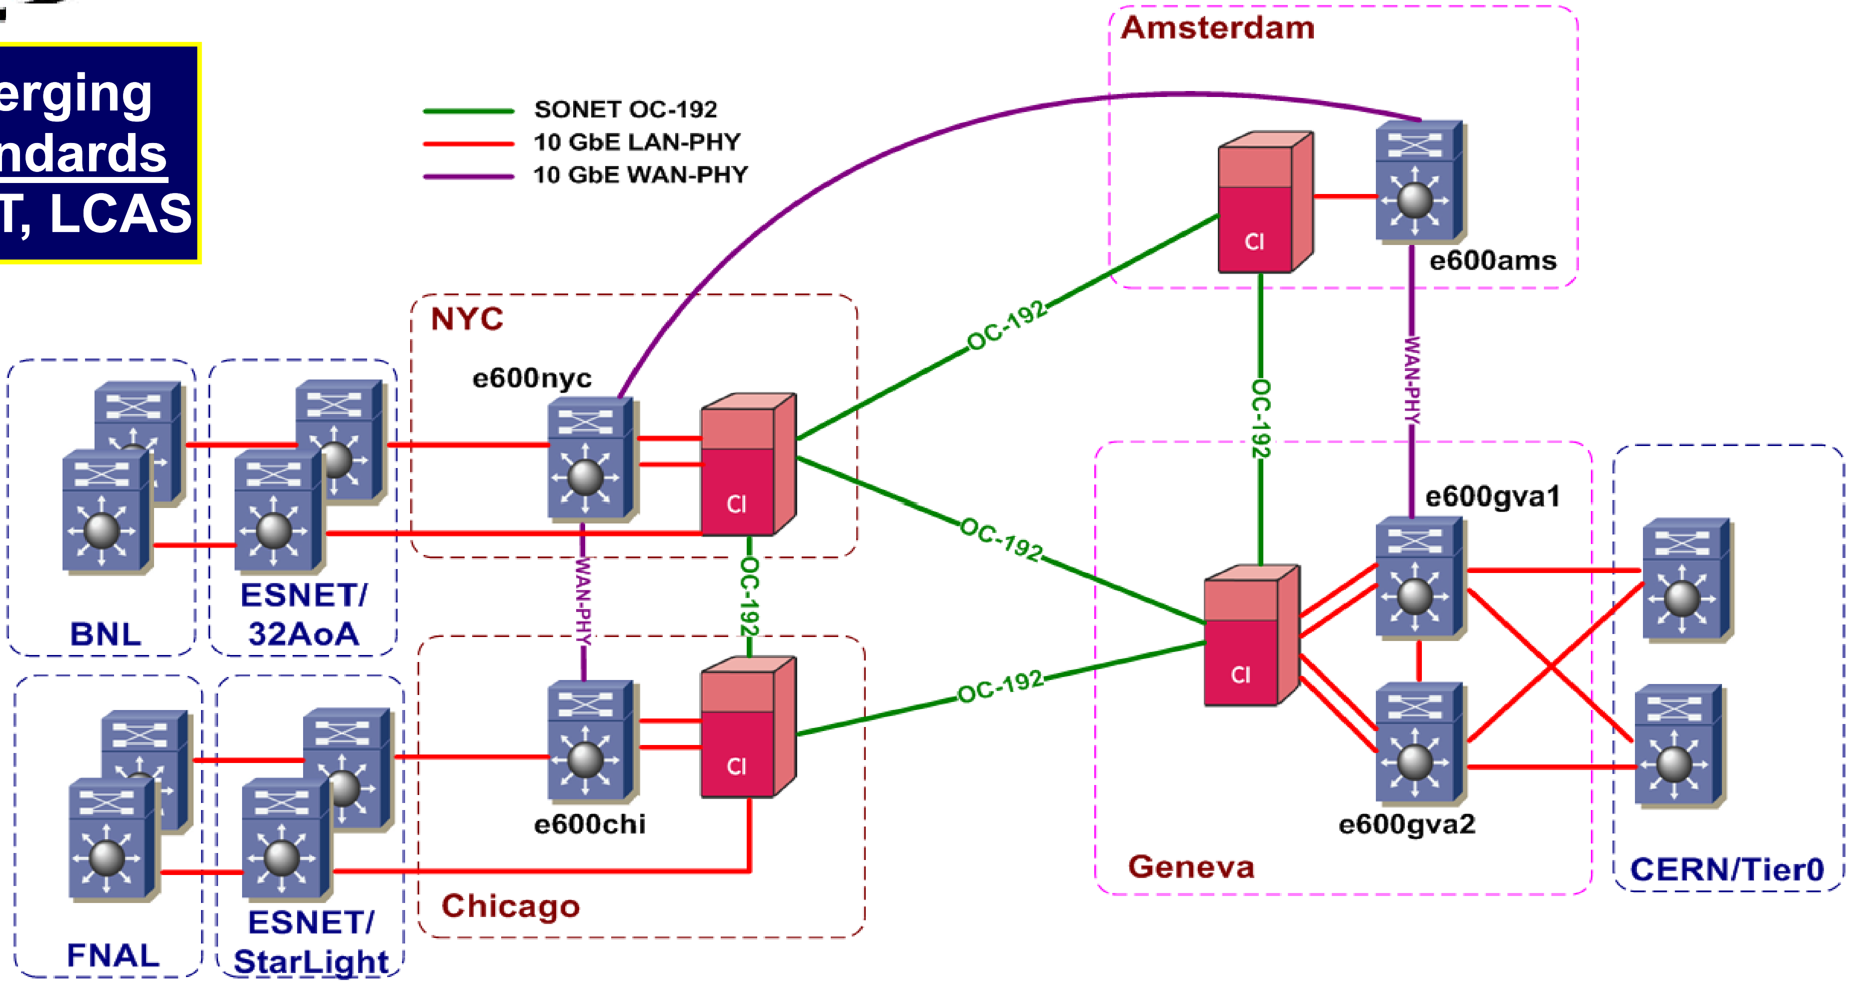

Robust fallback at layer 1 + next-generation hybrid optical network:
***Dynamic* circuit-oriented network services with BW guarantees**

Accommodation of Multi-layer Services

- L3 and L2 traffic are accommodated in shared bandwidth by L2 multiplexing and transferred to IP router, where each traffic is encapsulated by MPLS labels as needed.
- L1 traffic is assigned dedicated bandwidth and separated from L2/3 traffic.
- L2/3 (or IP/MPLS) traffic bandwidth can be hitlessly changed by LCAS to flexibly accommodate multi-layer services.

Japan

* Multi-protocol Label Switching (MPLS); Link Capacity Adjustment Scheme (LCAS)

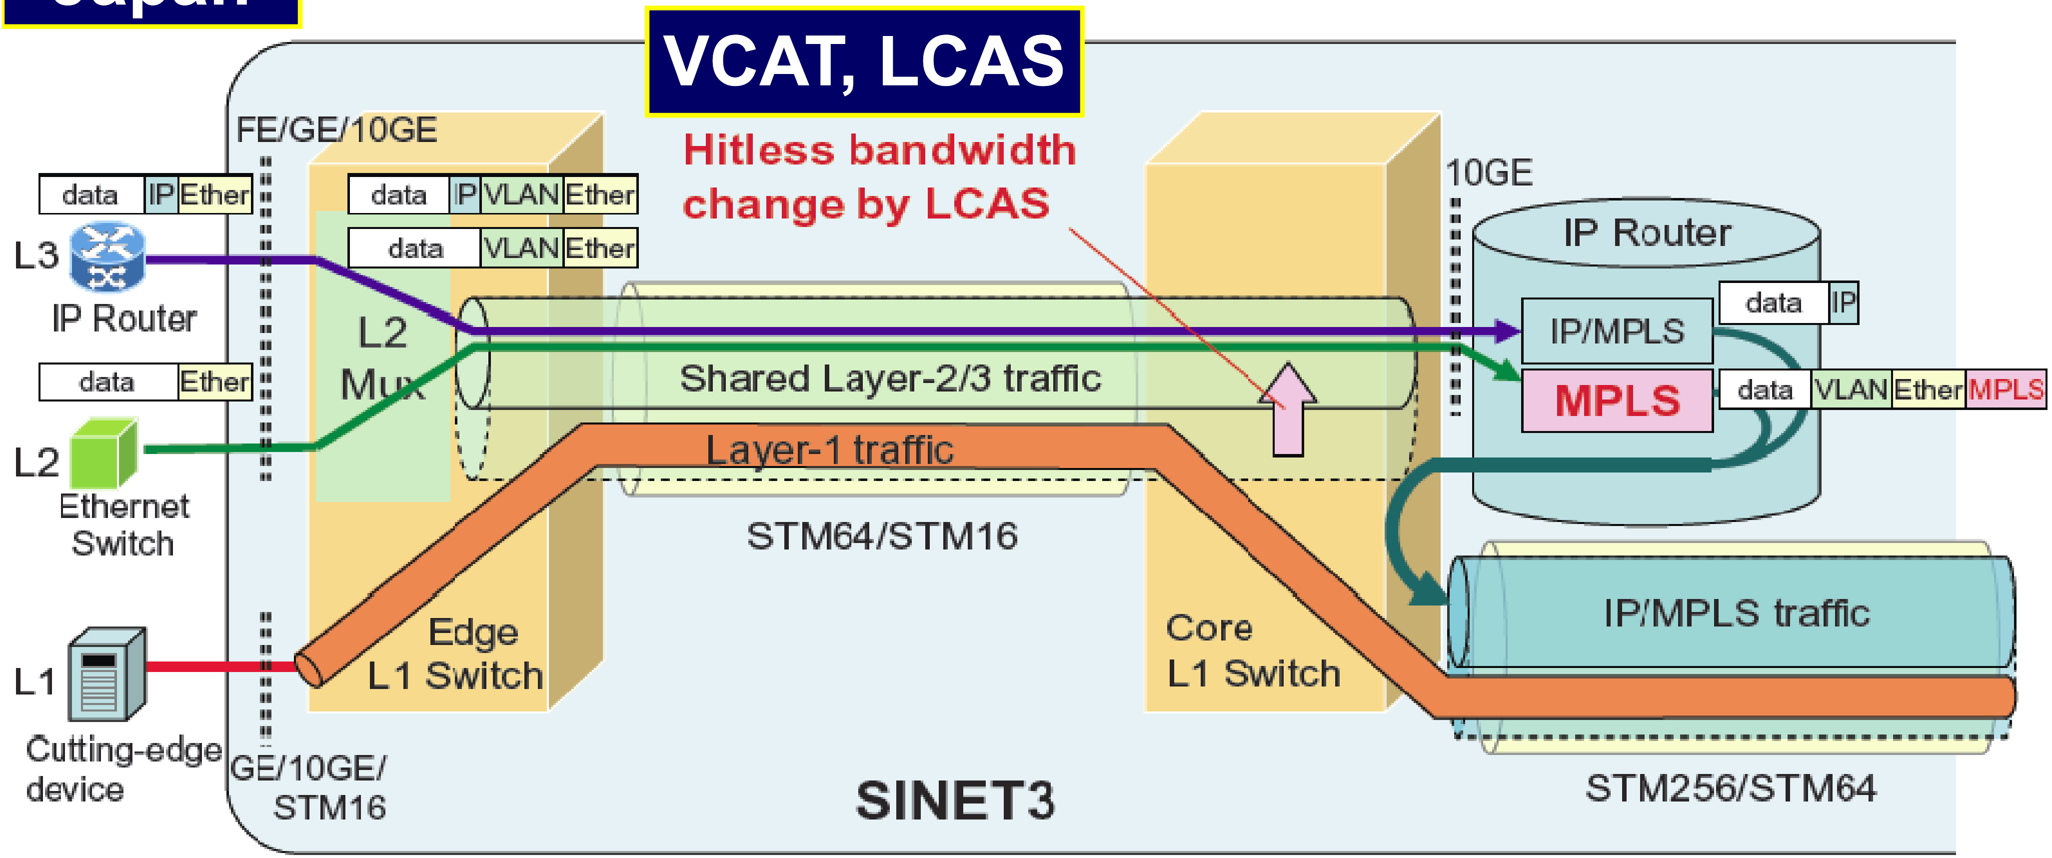

Network Services for Managed End-to-End Data Transfers

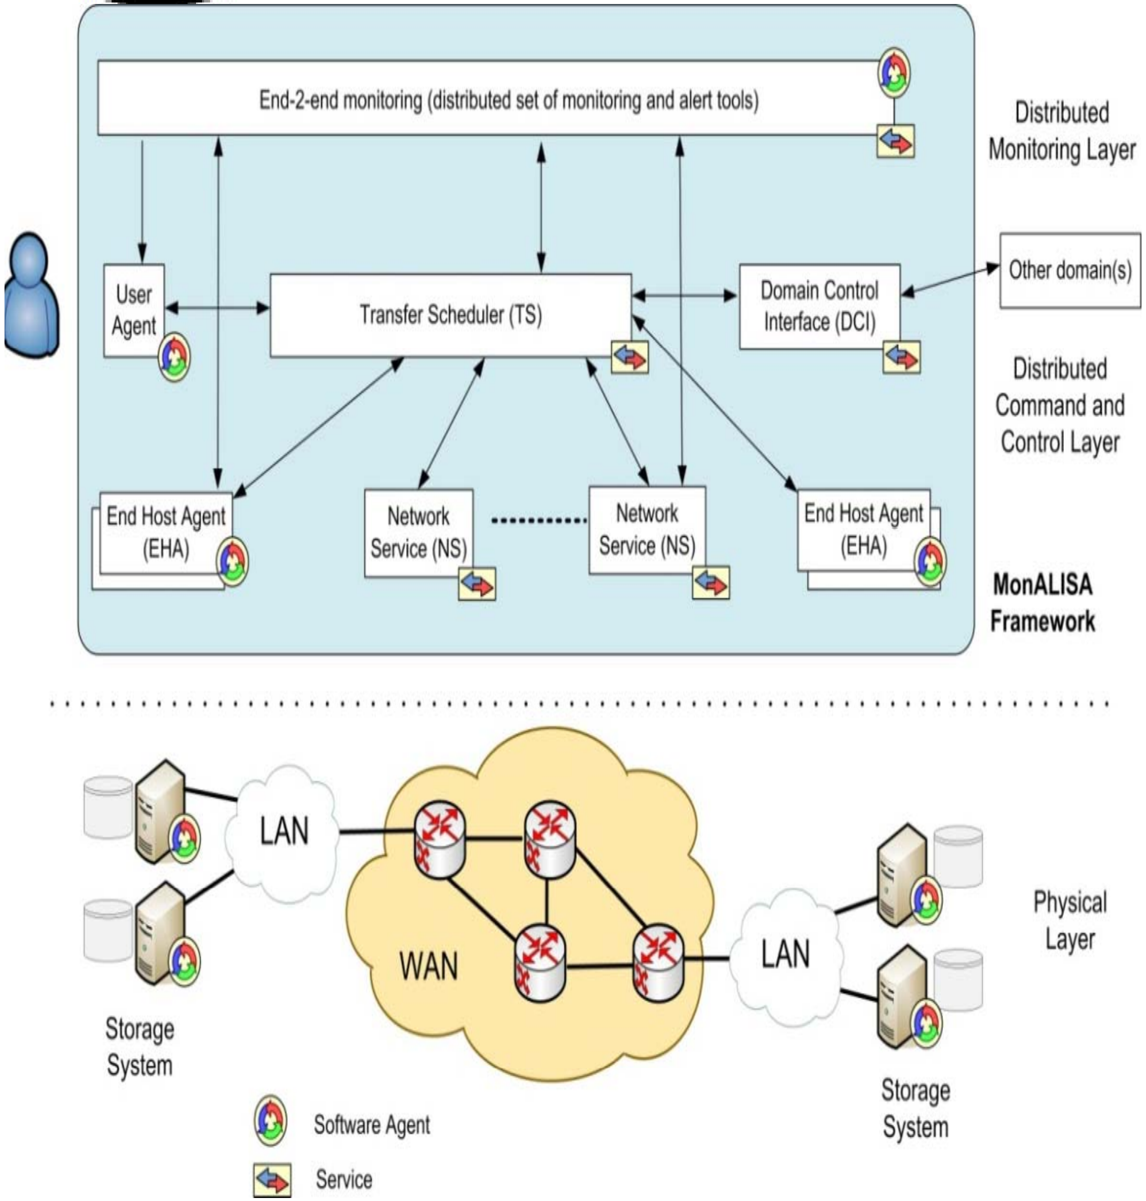

Robust Network Services based on

- Bandwidth guarantees
  - Virtual Circuits
- Scheduled Transfers
  - Transfer Classes
  - Priorities
- Monitoring of all components end-to-end
  - Network Elements
  - End-Hosts
- Interface to other circuit-oriented systems
  - Be part of heterogeneous end-to-end infrastructure

US LHC NWG

Problem Finding and Resolution

- Problems encountered today are hard to track due to missing the *global view* of the system

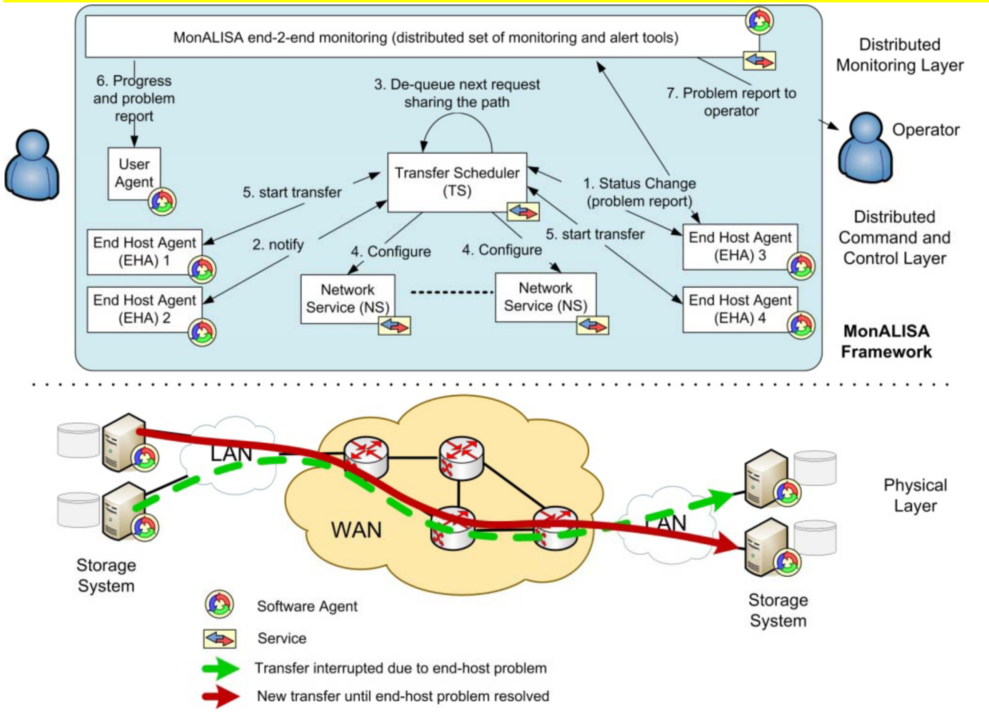

- **Example situation:** the system recognizes an end-host problem during the transfer and takes mitigating actions, re-scheduling transfers and notifying operators

End-to-end Monitored Managed Transfers

- Track problem to the source
  - Network / End-host
- Take appropriate action
  - Change transfer path
  - Adjust end-host parameters
  - Re-schedule transfer
- Provide experts with relevant (real-time) information
- Keep the user / application up-to-date on transfer progress
- Progessive automation: Target optimal resource utilization

- Developed in the field-proven MonALISA Framework

MonALISA: Monitoring Grids, Networks, Compute Nodes, Running Jobs, Processes

- Built for speed and global scale: 5k messages/sec/server; multi-threaded engine schedules ML services
- Autonomous agents auto-discover and collaborate in real-time for a variety of tasks

CENIC Innovation Award '06

Internet2 IDEA Award '07

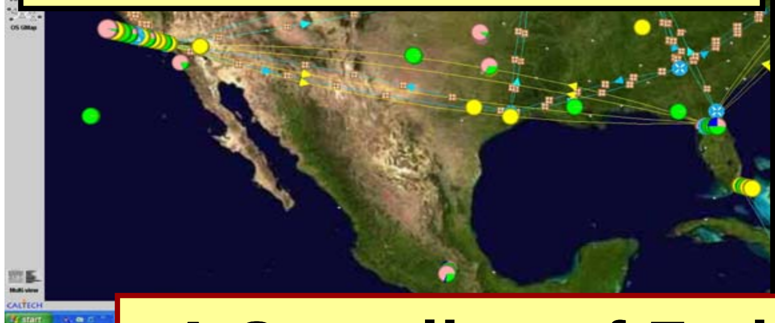

MonALISA Today

Running 24 X 7 (5 Years)
Now at 340 Sites

- Collecting > 1,000,000 parameters in near real-time
- Update rate of >20,000 parameter-updates per sec
- Monitoring
  - > 40,000 CPUs
  - > 100 WAN Links
- Many Thousands of Grid jobs running concurrently

Major Communities

- OSG
- CMS
- ALICE
- D0
- STAR
- VRVS, EVO
- LGC RUSSIA
- SE Europe GRID
- APAC Grid
- UNAM Grid (Mx)
- ITU

- ABILENE
- ULTRALIGHT
- GLORIAD
- US LHCNet
- RoEduNET
- Enlightened

A Sampling of End To End Monitoring and Global views

Dynamic Network Path Allocation & Automated *Dataset* Transfer with Optical Switches

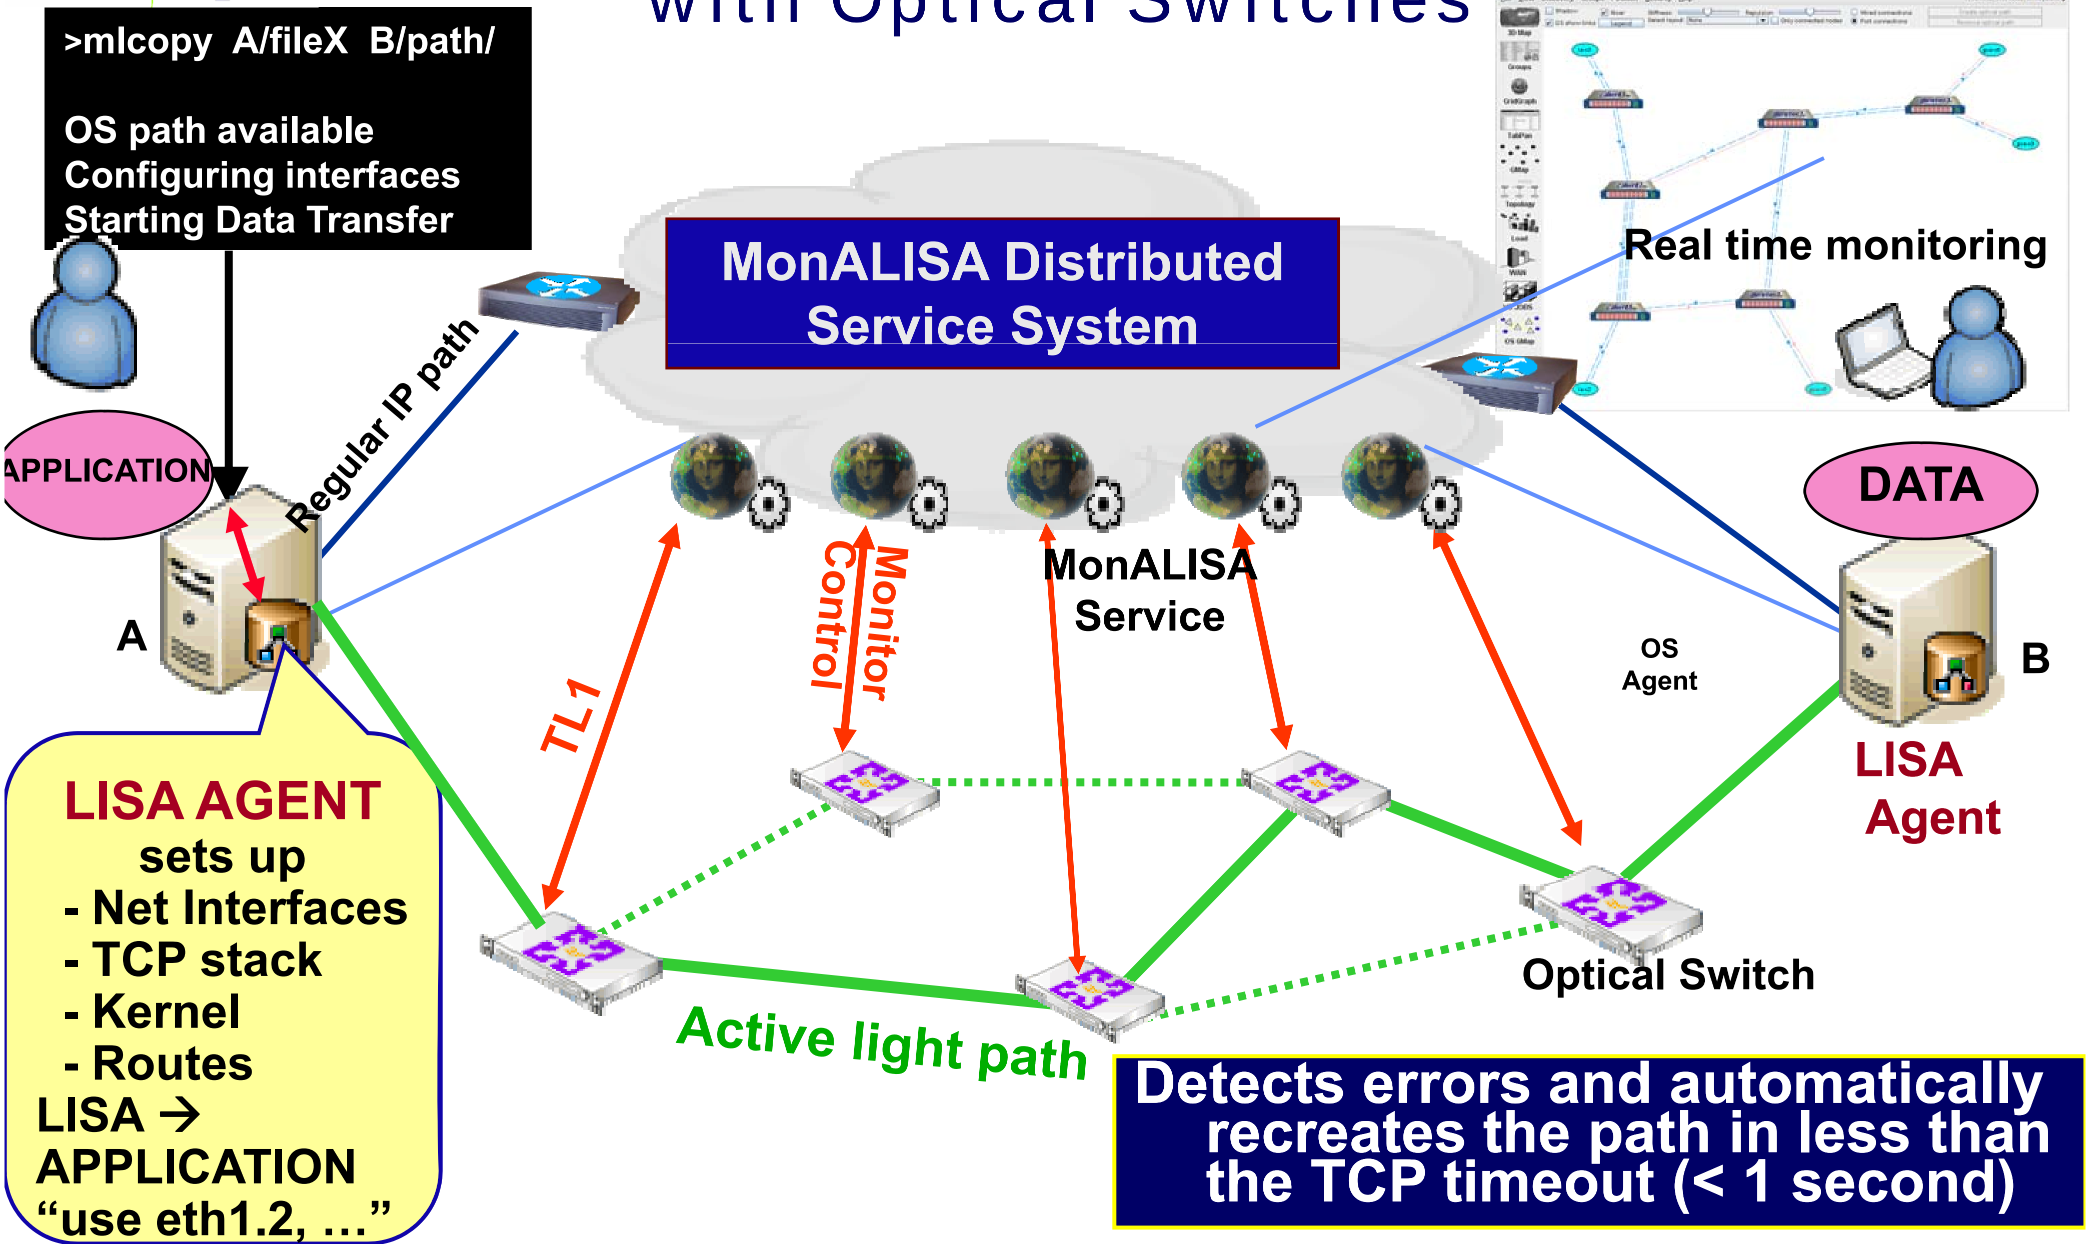

FDT Automatic Path Recovery: Fiber Cut Simulations

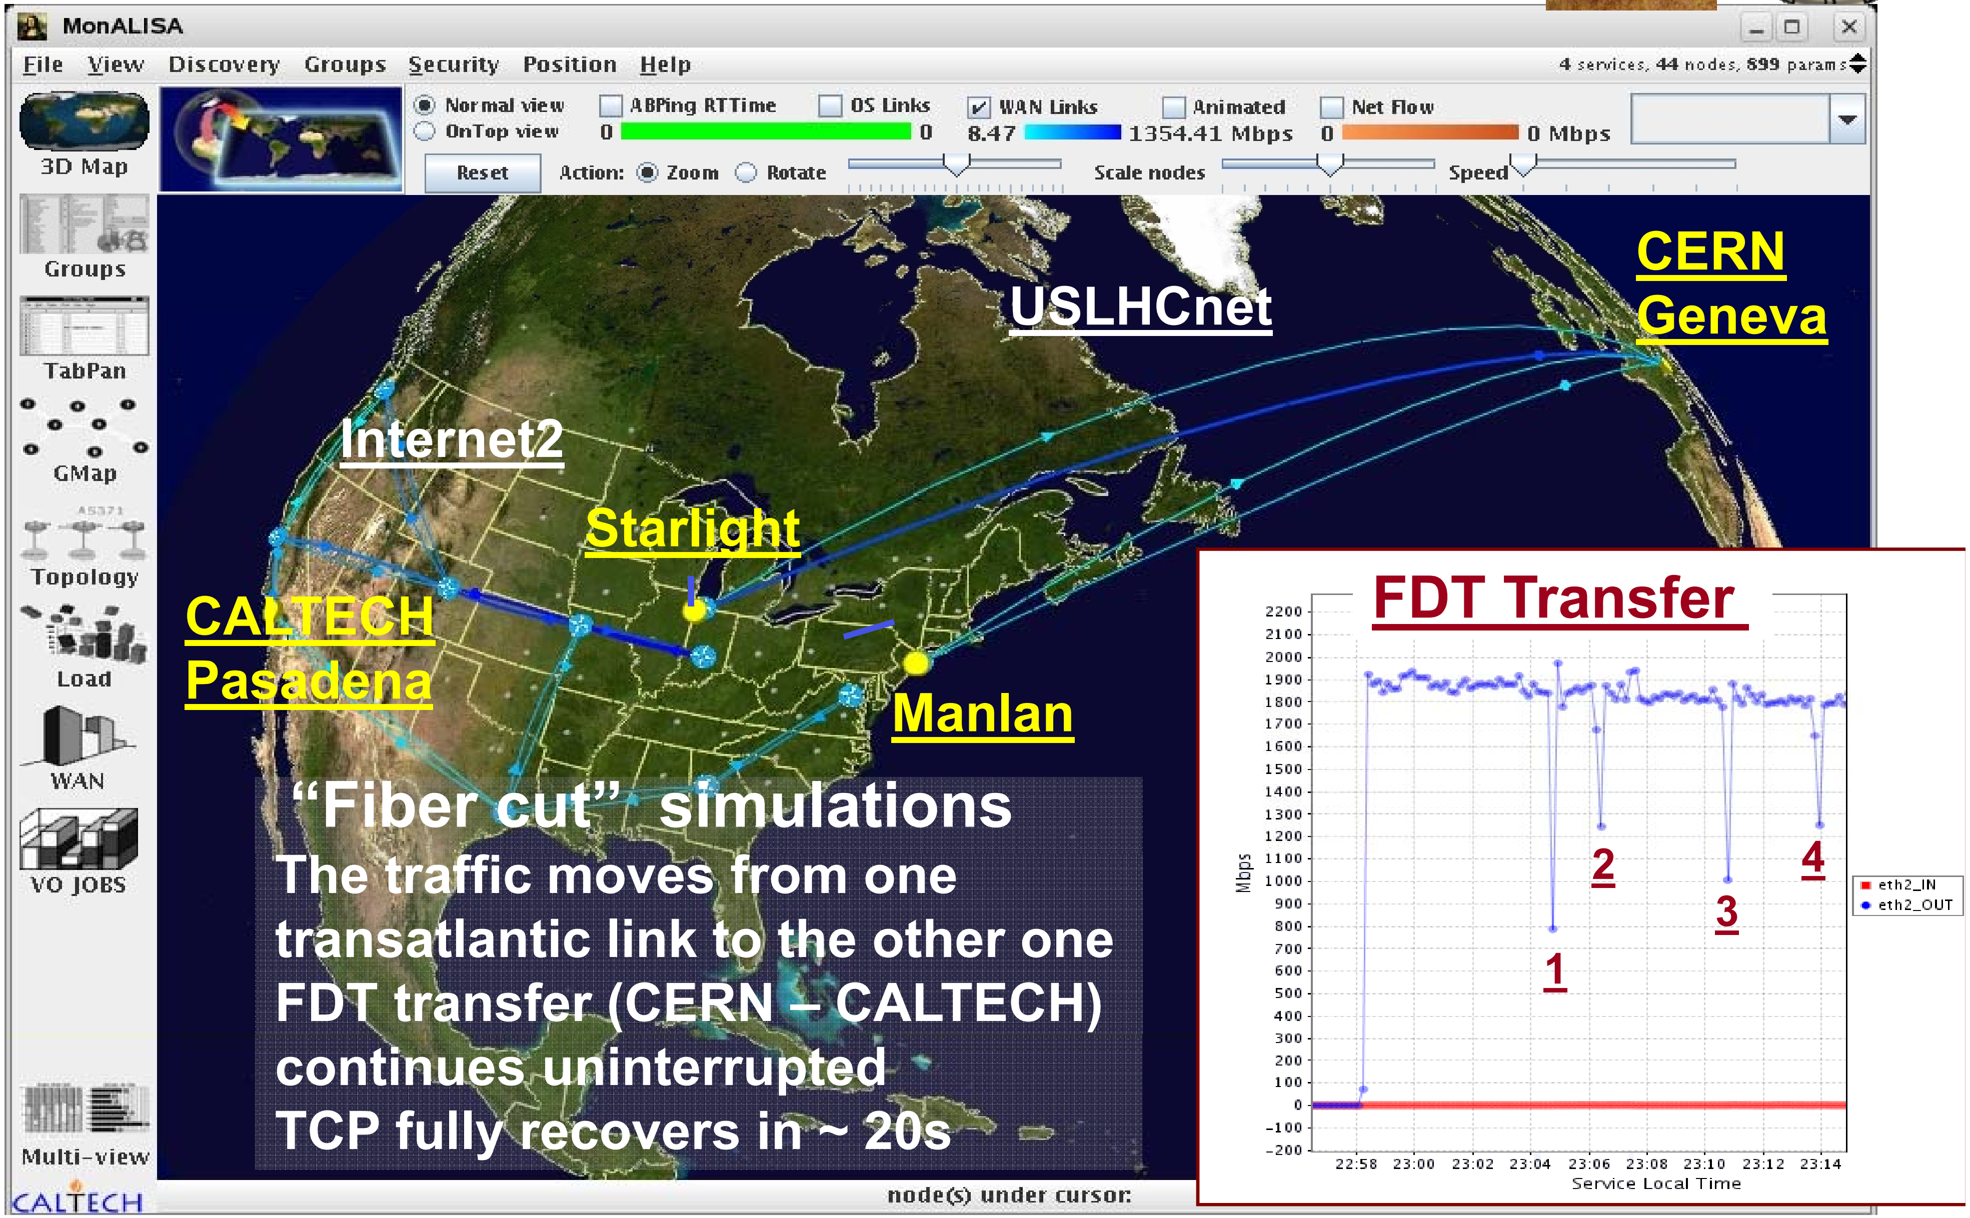

The Main VINCI Services

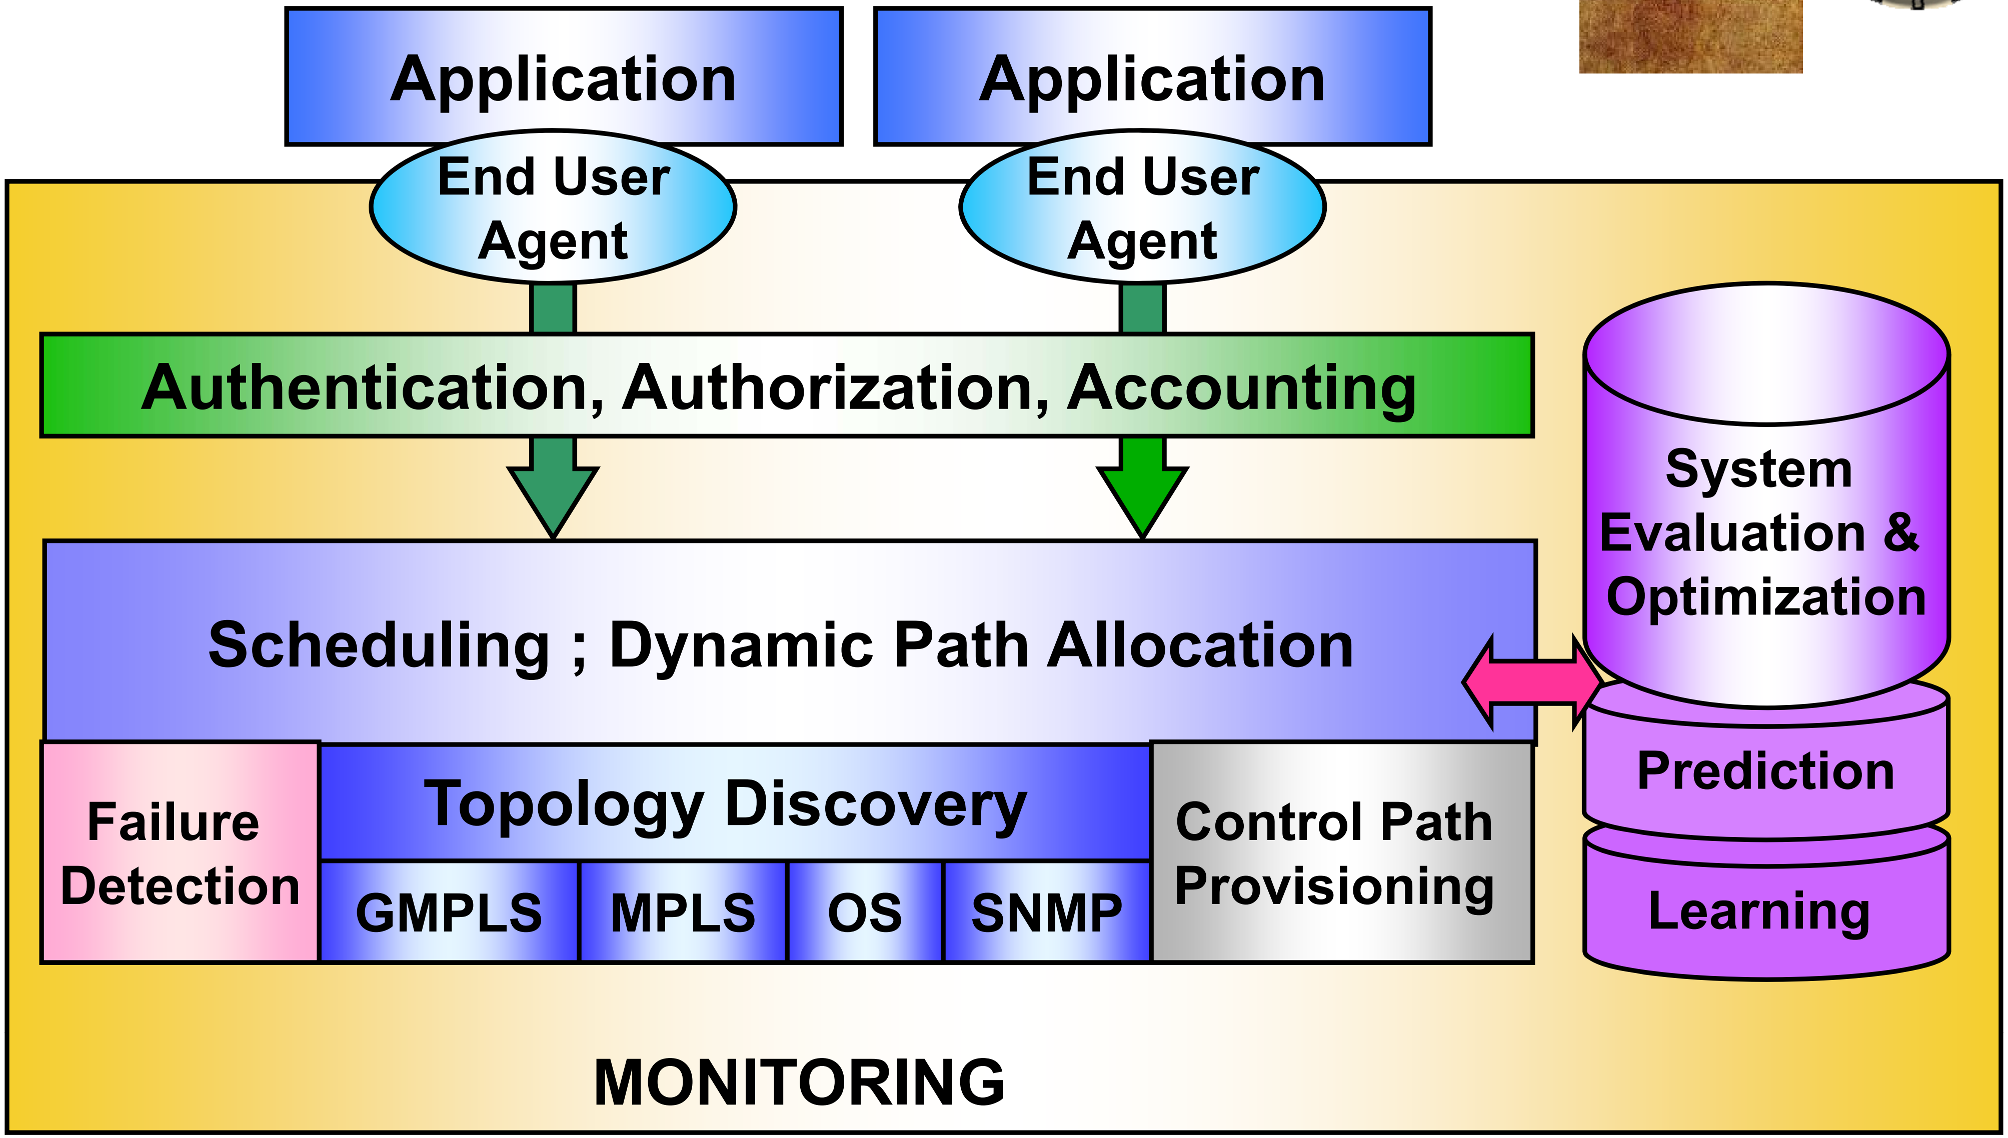

SCIC Main Conclusion for 2007

- ***As we progress we are in danger of leaving the communities in the less-favored regions of the world behind***
- ***We must Work to Close the Digital Divide***
  - ➔ ***To make physicists from all world regions full partners in the scientific discoveries***
  - ➔ ***This is essential for the health of our global collaborations, for our field, and for the world community***

Work on the Digital Divide from Several Perspectives

- **Share Information: *Monitoring,* Tracking *BW Progress*; Dark Fiber Projects & Pricing**
  - **Track Planning (focus on LHC) and Leading Edge Progress**
  - **Model Cases: Brazil, Poland, Slovakia, Czech Rep., China …**
  - **Encourage Access to Dark Fiber; Modern technology choices**
- **Raise Awareness: Locally, Regionally & Globally**
  - **Digital Divide Workshops**
  - **Diplomatic Events: WSIS, RSIS; Bilateral: e.g. US-India**
- **Technical Help with Modernizing the Infrastructure:**
  - **Provide Tools for Effective Use: Data Transport, Monitoring, Collaboration**
  - **Design, Commissioning, Development**
- **Encourage, and Work on Inter-Regional Projects**
  - **GLORIAD, Russia-China-Korea-US-Europe Optical Ring**
  - **Latin America: CHEPREO/WHREN (US-Brazil); RedCLARA**
  - **Mediterranean: EUMEDConnect; Asia-Pacific: TEIN2**

SCIC Monitoring WG
PingER (Also IEPM-BW)

R. Cottrell

Monitoring & Remote Sites (1/07)

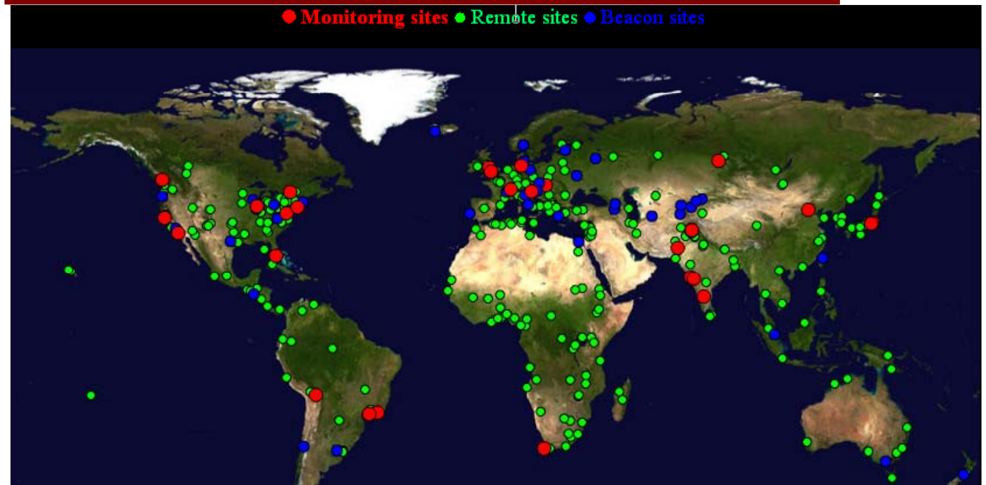

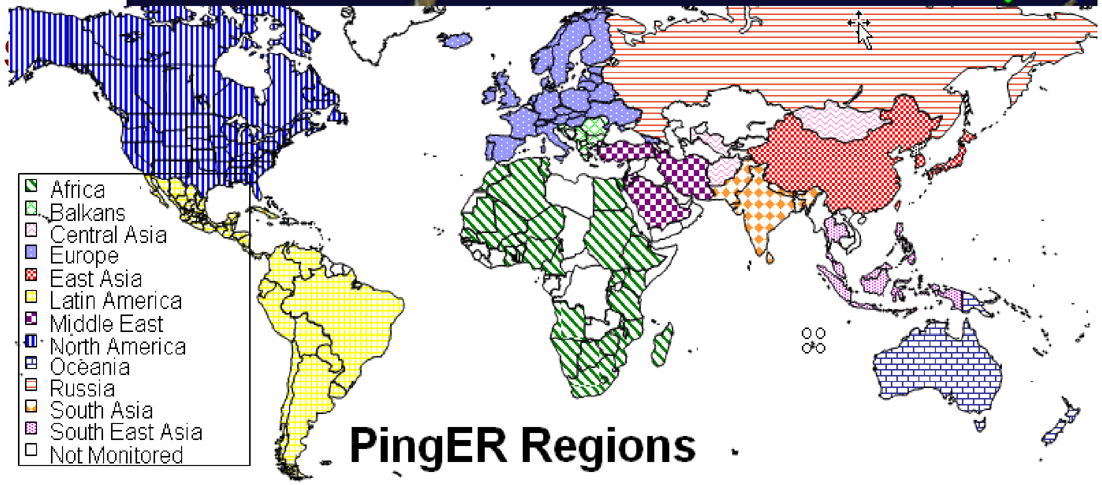

- **Measurements from 1995 On**
  - ***Reports link reliability & quality***
- **Countries monitored**
  - ➔ **Contain 89% of world pop.**
  - ➔ **99% of Internet users**
- **600+ remote sites in 115 nations**
  87 Sites in 32 African countries
- **30 monitors in 14 countries**
  Capetown, Rawalpindi, Bangalore
- **Excellent Work; funding issue**

Countries: N. America (2), Latin America (16), Europe (22), SE Europe (6), Africa (32), Mid East (8), Central Asia (9), South Asia (5), East Asia (4), SE Asia (6), Russia (1), China (1) and Oceania (4)

SCIC Monitoring WG – Throughput Improvements 1995-2007

Progress: but Digital Divide is Mostly Maintained

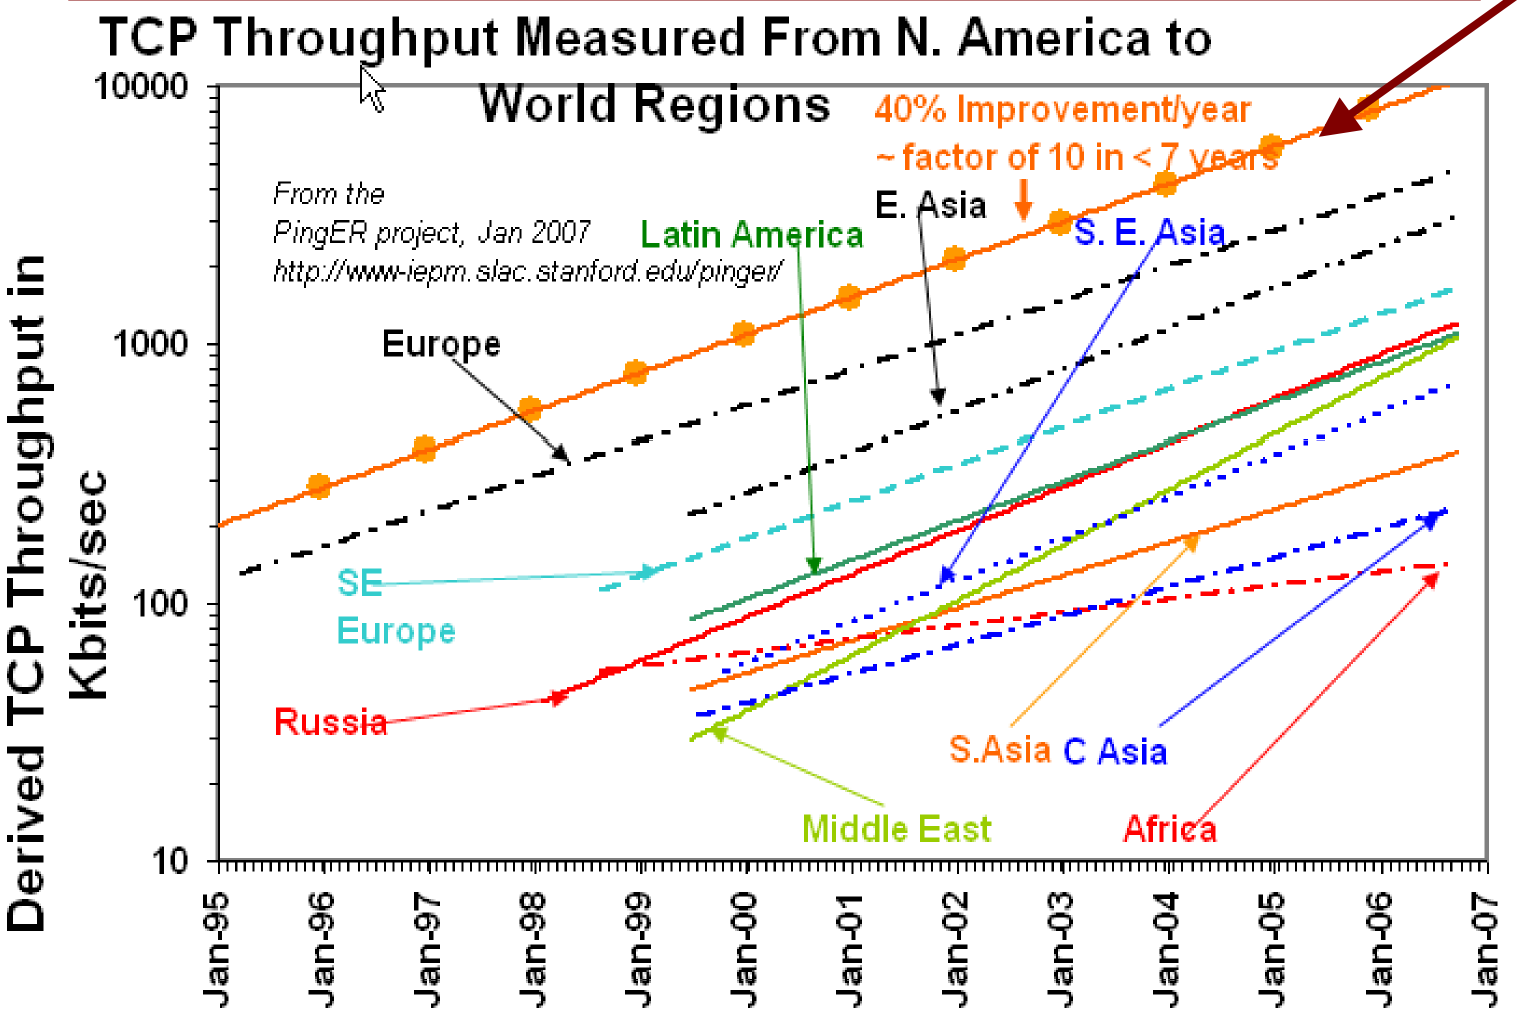

40% annual improvement ~10X Per 7 yrs

Behind Europe
- **6 Yrs: Russia, Latin America**
- **7 Yrs: Mid-East, SE Asia**
- **8-9 Yrs: So. Asia**
- **11 Yrs: Cent. Asia**
- **12 Yrs: Africa**

India, Pakistan, Central Asia, and Africa are in Danger of Falling Even Farther Behind

R. Cottrell

Bandwidth of TCP < MSS/(RTT*Sqrt(Loss))

Matthis et al., Computer Communication Review 27(3), July 1997

Throughput from US to African Countries vs UN Human Development Index

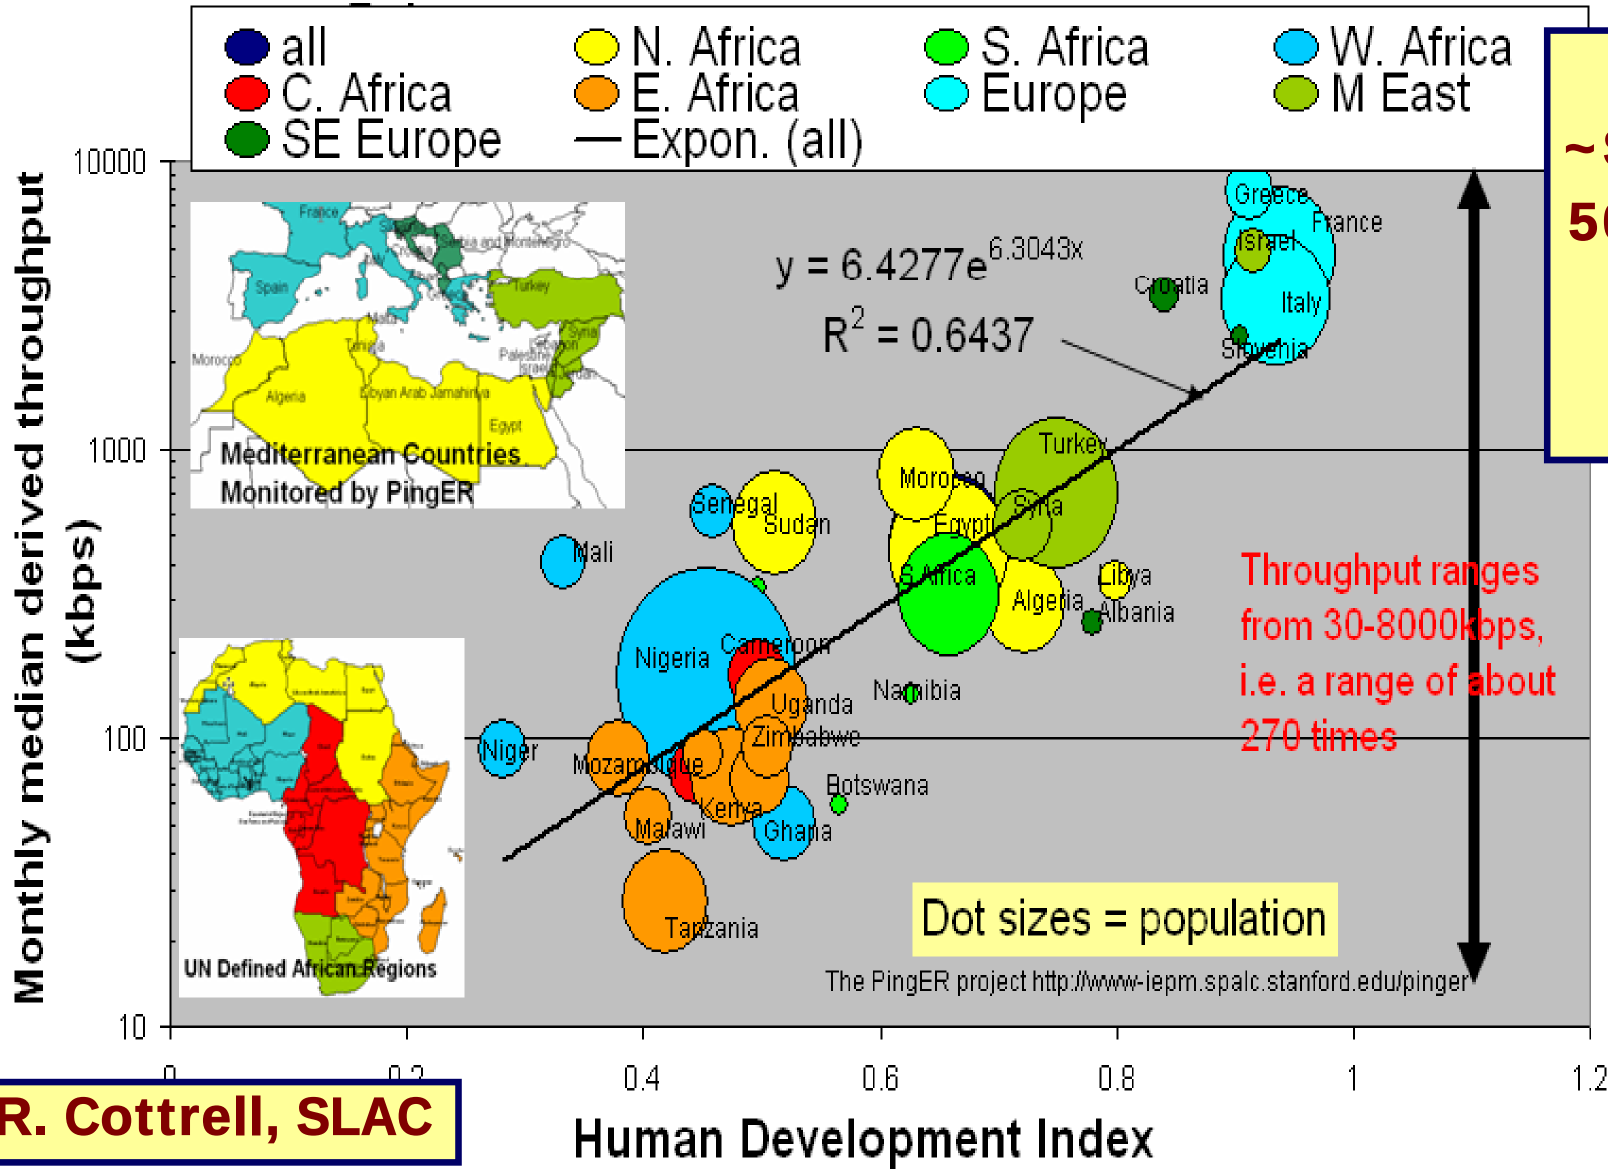

COSTS

~$2.5/kbps/Mo

50-1000X More than well-developed regions

R. Cottrell, SLAC

SCIC Digital Divide Workshops and Panels

- ***2002-2007:***
  An effective way to raise awareness of problems, and discuss approaches and opportunities for sol'ns with national and regional communities, and gov't officials
  - ***ICFA Digital Divide Workshops: Rio 2/2004; Daegu 5/2005***
  - ***CERN & I2 Workshops on R&E Networks in Africa***
  - ***World Summit on the Information Society,***
- ***In 2006***
- ***February: CHEP06 Mumbai: Digital Divide Panel, Network Demos, & Workshop [SCIC, TIFR, CDAC, Internet2, Caltech]***
  "Moving India into the Global Community Through Advanced Networking"
- ***May-June: US-India Summit, NAS Meeting, Meeting at Cal State LA***
- ***October: ICFA Digital Divide Workshops in Cracow & Sinaia***
- ***April 14-17 2007 at APS: "Bridging the Digital Divide" Sessions***
  Sponsored by Forum for International Physics

Brazil: RNP2 Next-Generation Backbone

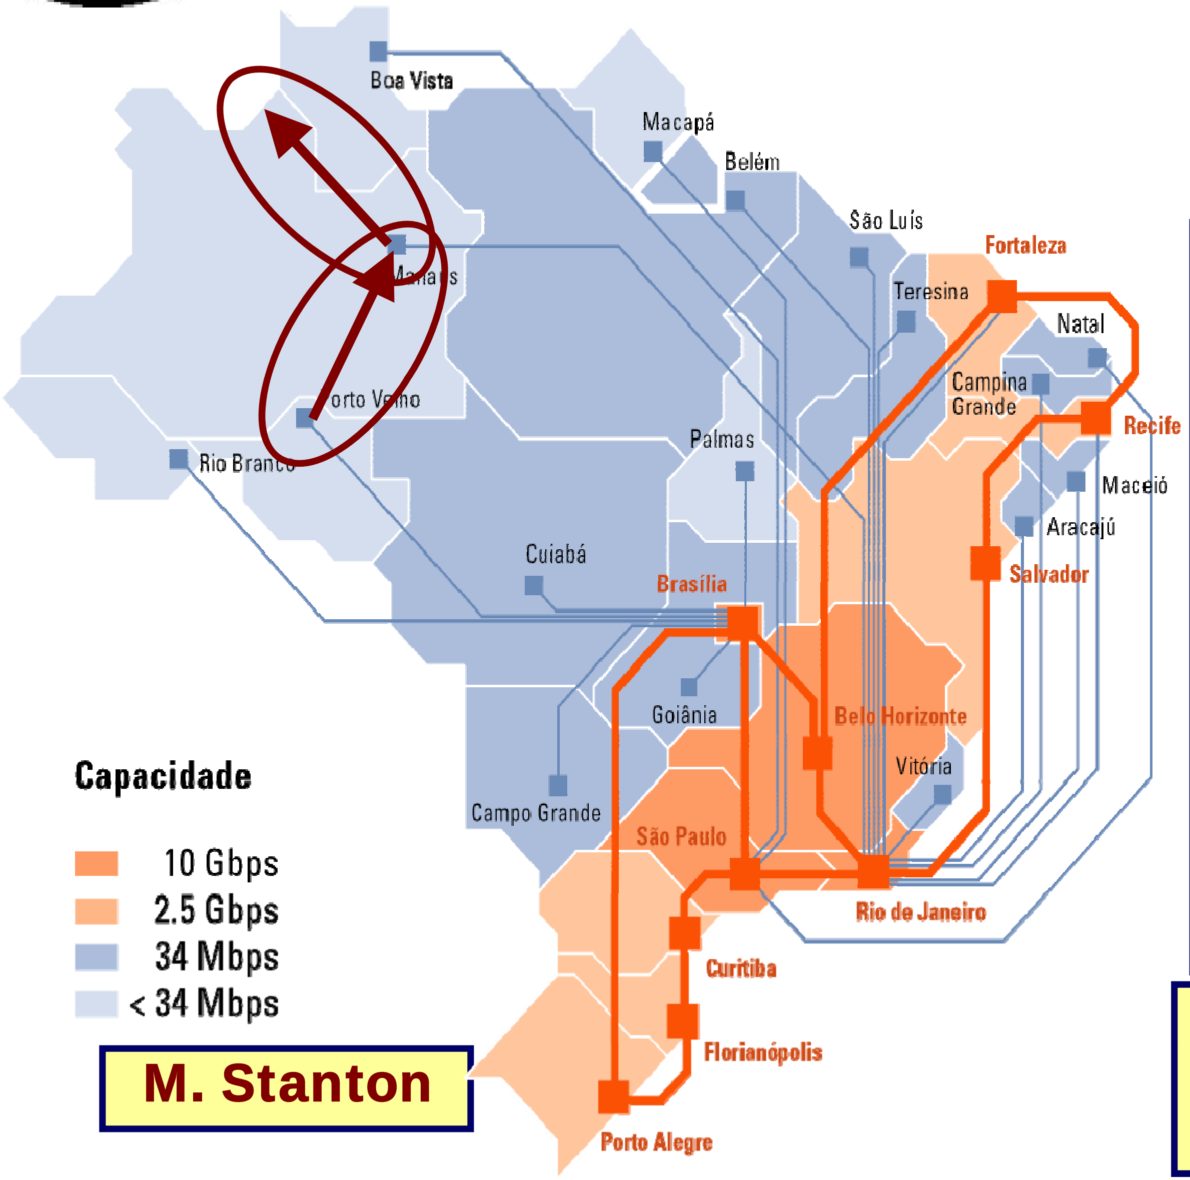

M. Stanton

New vs. Old
A factor of 70 to 300 in Bandwidth

2006-7:

- **Buildout of dark fiber nets in 27 cities with RNP PoPs underway**
- **200 Institutions Connected at 1 GbE by End-2008 (Well-advanced)**
- **2.5G (to 10G) WHREN (NSF/RNP) Link to US; 622M Link to GEANT**

Now extending to the Northwest; Dark fiber across the Amazon to Manaus

Optical Fiber Through the Amazon: Porto Velho-Manaus

935 Km Optical Cables
8,000 Stakes
7 Optical Repeaters
(622 Mbps)
60m Towers
Microwave across
the Amazon

CD

1 GbE Connections to UERJ and UNESP Centers: RNP and the GIGA Project

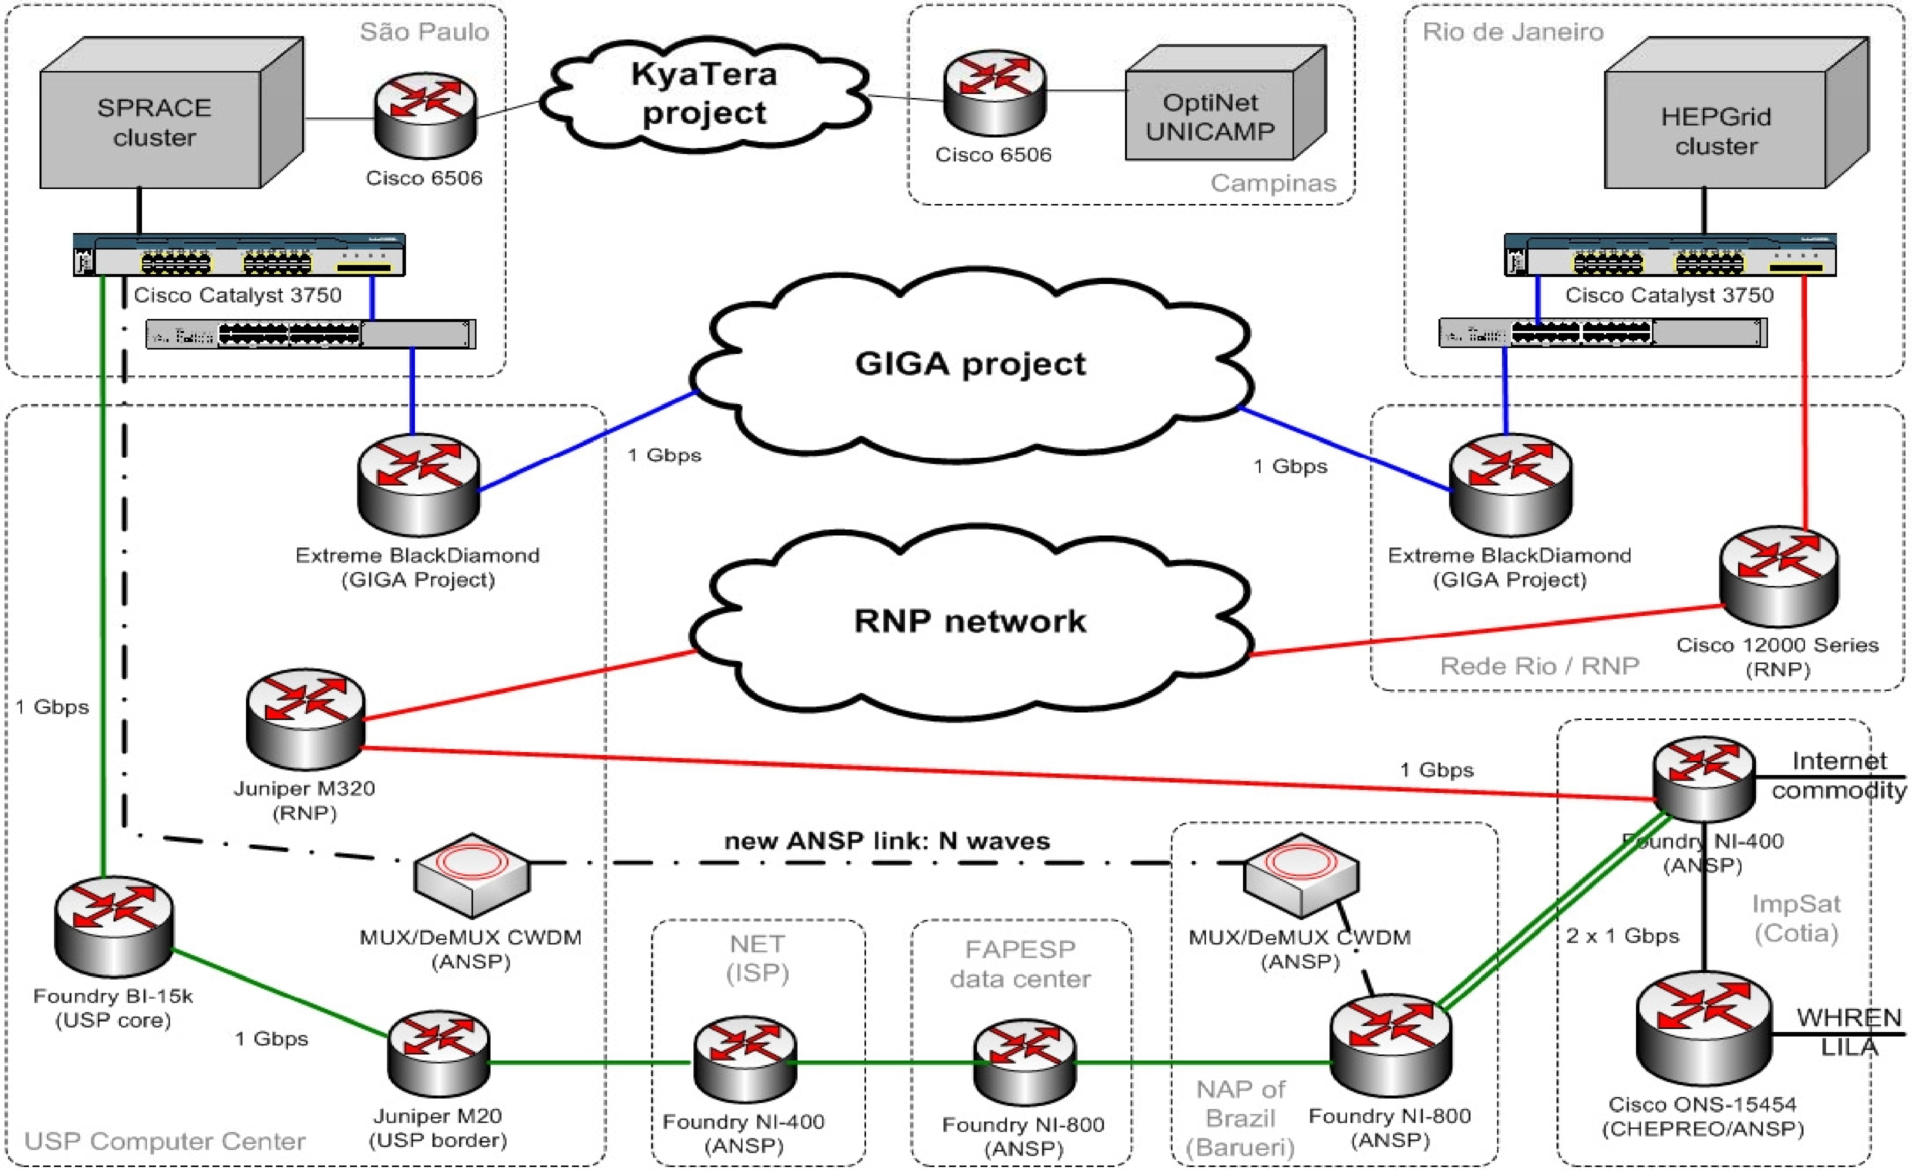

Mumbai-Japan-US Links

INDIA
Chennai POP VSNL
LANDING STATIONS

VSNL MUX
TIC MUX

SINGAPORE LANDING STATION

TIC
Cable

JAPAN
JAPAN LAND STANDING

TIFR Link to Japan
+ Onward to US & Europe

Loaned Link from VSNL at CHEP06

End to End Bandwidth
4 X 155 Mbps
on SeMeWe3 Cable

Goal is to Move to
10 Gbps on SeMeWe4

Helped spark planning for Next Generation R&E "Knowledge Network" in India

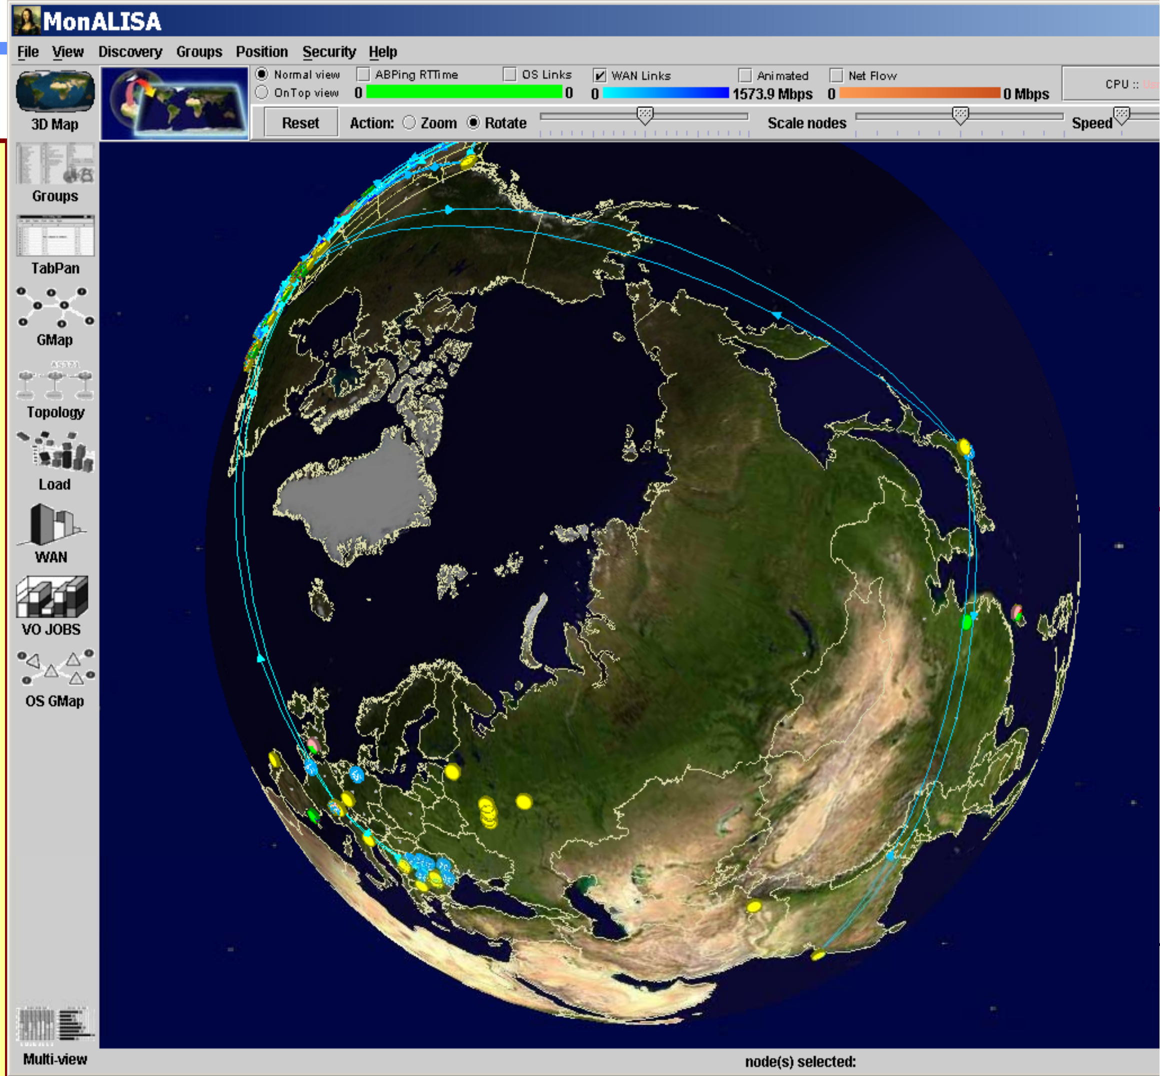

INTERFACE TYPES
STM 4

INTERFACE TYPES
OC-12

TIFR Mumbai, INDIA

NTT Otemachi Bldg, JAPAN
+ Onward to US, Europe ➡

President of India Kalam Collaborating with US, CERN, Slovakia via VRVS/EVO

Coincident with Data Transfers of ~500 Mbps

15 TBytes to/from India in 2 Days

Helped spark planning for Next Generation "Knowledge Network" in India

India: Knowledge Commission Recommendation to Create a National "Knowledge Network": Approved by Prime Minister January 25, 2007

- *"Build a National Knowledge Network with gigabit capabilities to connect all universities, libraries, laboratories, hospitals and agricultural institutions to share data and resources across the country." [~ $2B]*
- *5000 Institutions; 500-1000 in Phase 1*
  [Time estimate: 3-6 Months]
- *Minimum connectivity at end nodes, 100 Mbps (to gigabit)*
- *Phase 1: Start with existing commercial networks*
  ➔ *"Slide" into hybrid network with inner core owned by the stakeholders*
- *Migrate core to N X 10 Gbps, providing gigabit connectivity*

Ongoing issue: 10 G int'l Links (affordable) to the Region [also Pakistan]

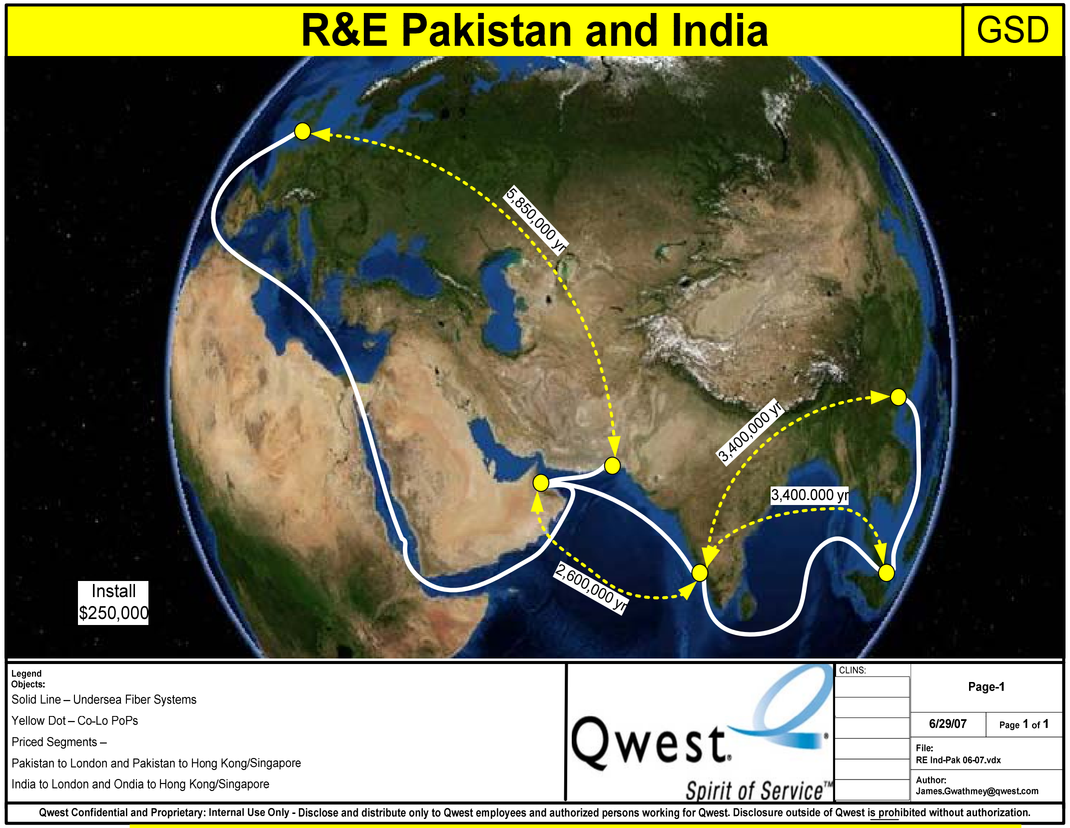

Pricing Issue: $ 3.4 - 8.5M Per Year Quoted for 622 Mbps
> 100X Cost/BW compared to Transatlantic Links

CERN Initiative by David Foster, CERN/IT

- Visited Reliance Telecom Mumbai
- 1 Gbps link Mumbai – CERN (TIFR Tier2)
- Full BW 6 Hrs. Per Day
- 300 Mbps available 24 hrs/day
- Better pricing than 45 Mbps link in place now

PERN2 Long Haul Fiber Layout Map

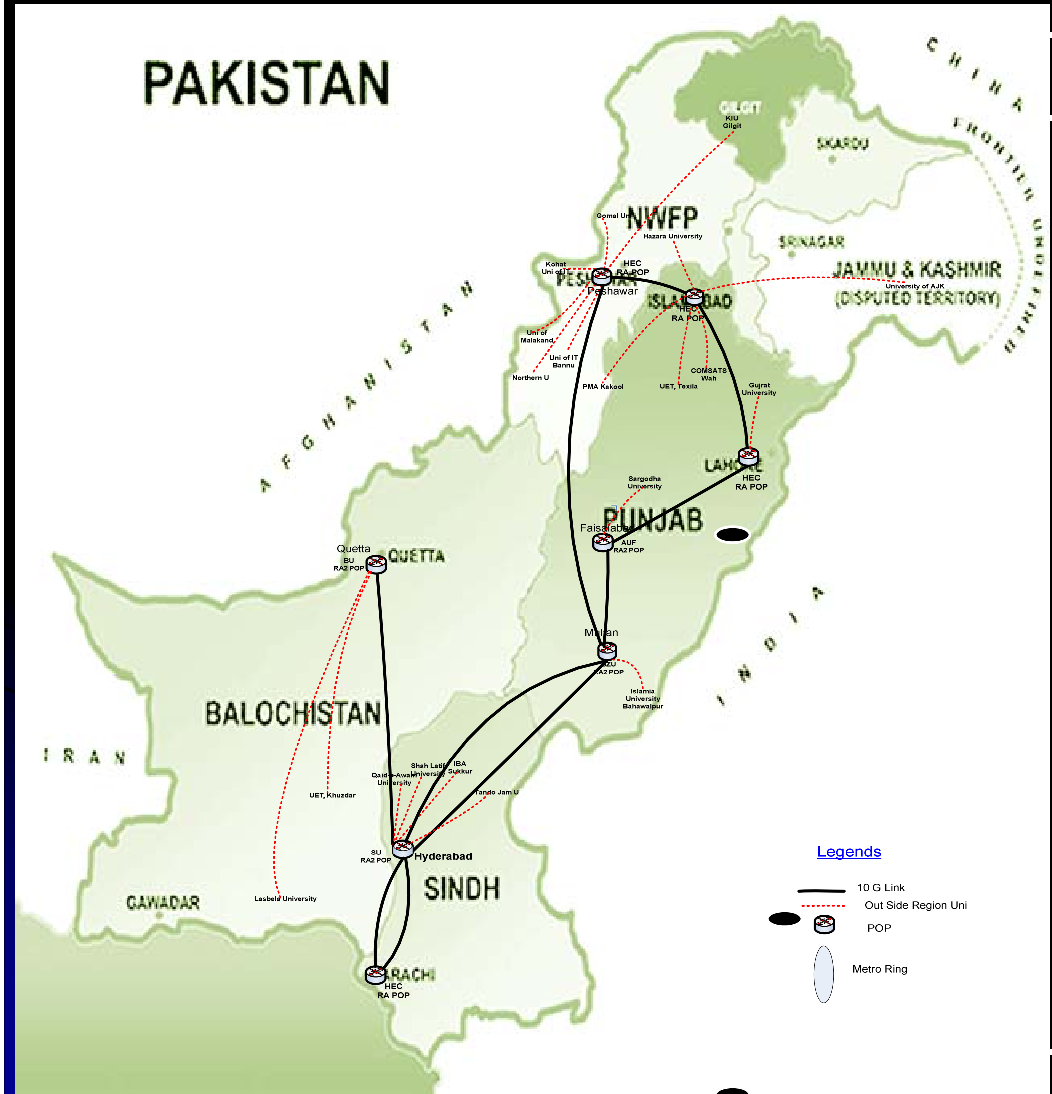

PERN: Next Three Years of Development (2005-08)

	2005	2007	By 2008
Institutes	59	85	200+
Core (Gbps)	0.050	0.155	10 [200x]
Last Mile (Gbps)	0.002	0.008	1 [500x]
No. of PoPs	3	3+2	7+3
Comm. Internet	0.155	2 x 0.155	0.622
NRENs Connec-tivity	Nil	Nil	*2.5 – 10 Gbps*

From 40 to 2000 Cities on Optical Fibre (2000-2007)

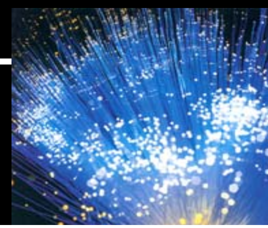

International ICFA Workshop on HEP Networking, Grid and Digital Divide Issues for Global E-Science

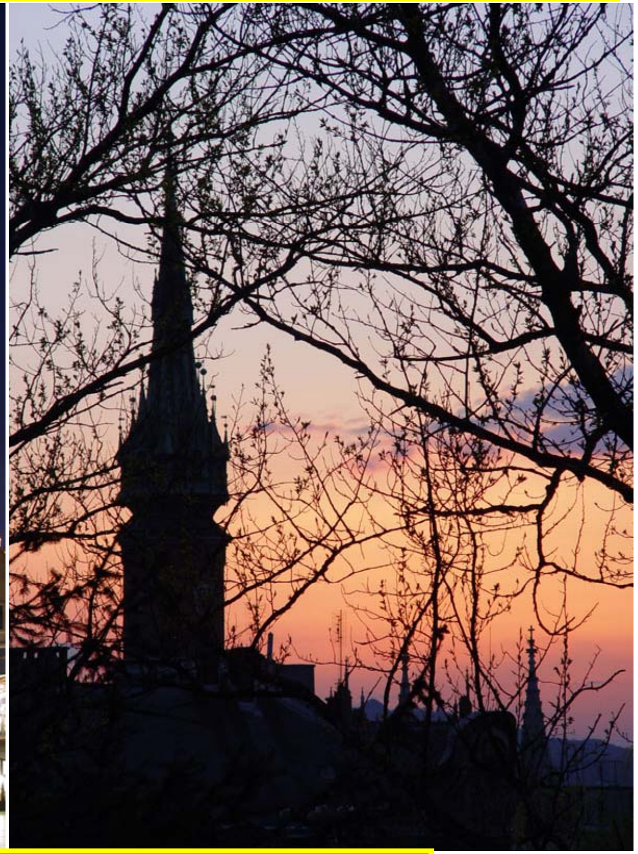

National Academy of Arts and Sciences
Cracow, October 9-11, 2006

http://icfaddw06.ifj.edu.pl/index.html

Michal Turala

PIONIER - THE POLISH OPTICAL INTERNET

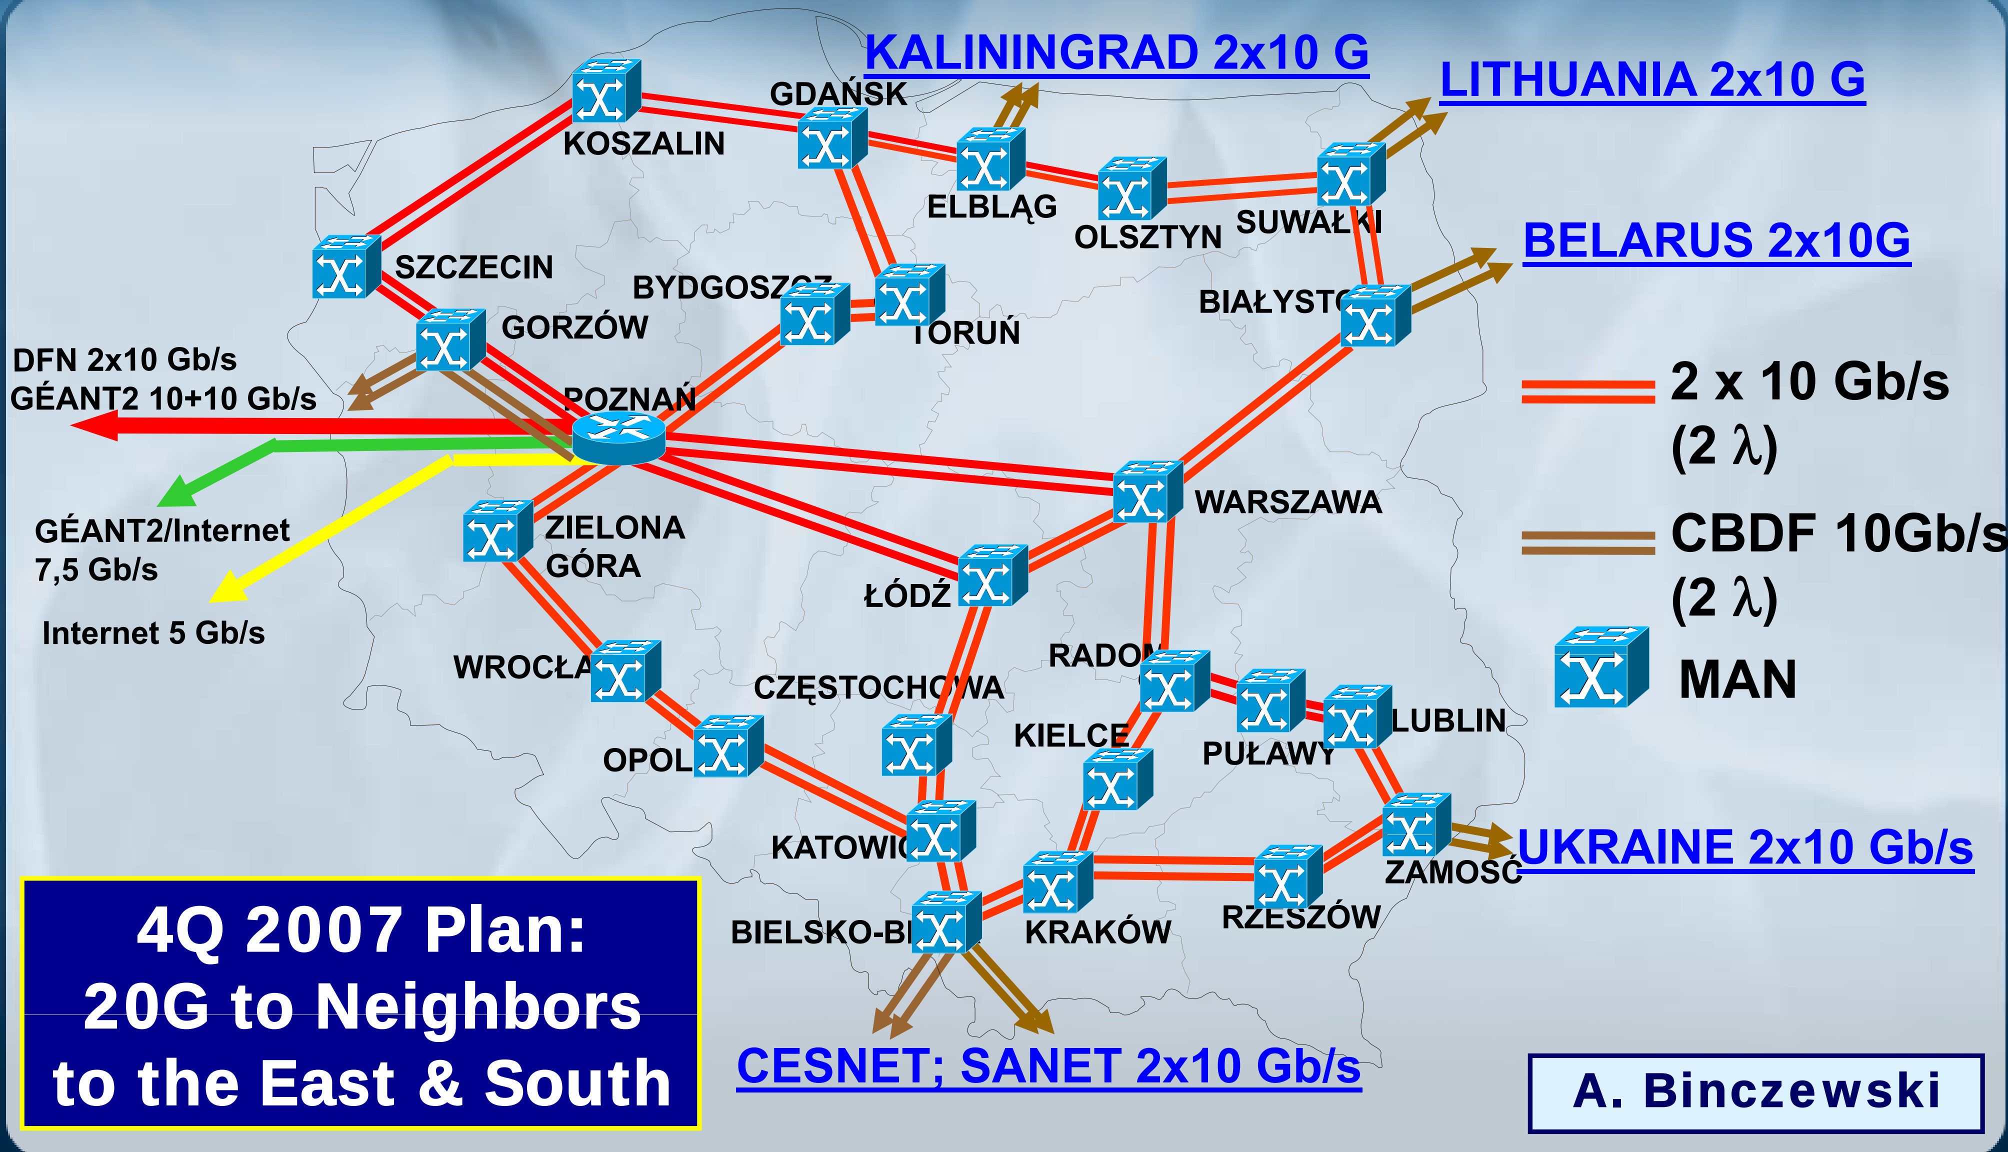

Cross Border Dark Fiber in Central Europe (2007)

Dark Fiber Rings:
Poland
Czech Republic
Austria
Slovakia

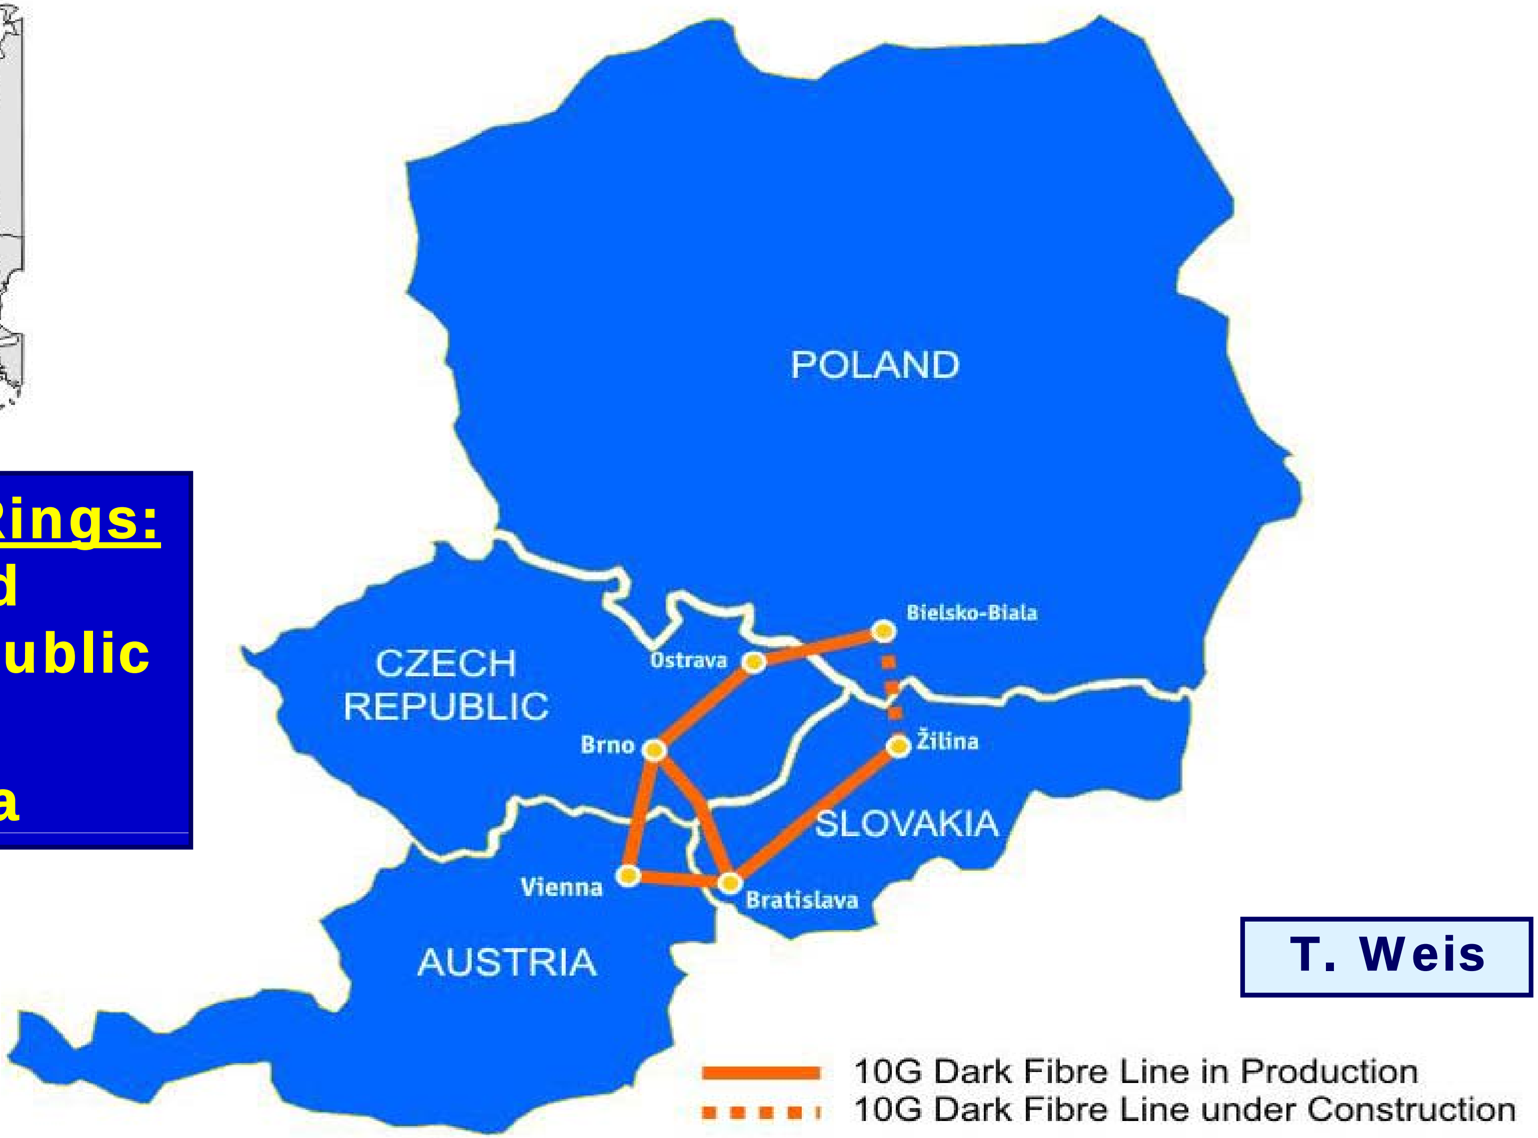

T. Weis

International ICFA Workshop on Grid Activities within Large Scale International Collaborations

Sinaia, Romania
October 13-18, 2006
http://niham.nipne.ro/events2006/

Mihai Petrovici

N. Tapus

Romania: Roedunet Topology Moving from 2-155 Mbps to 1 and then 10 Gigabit Ethernet in 2007

Q1 2007: Gigabit Ethernet Between Cities

End 2007: 10G Among Major Cities and Internationally

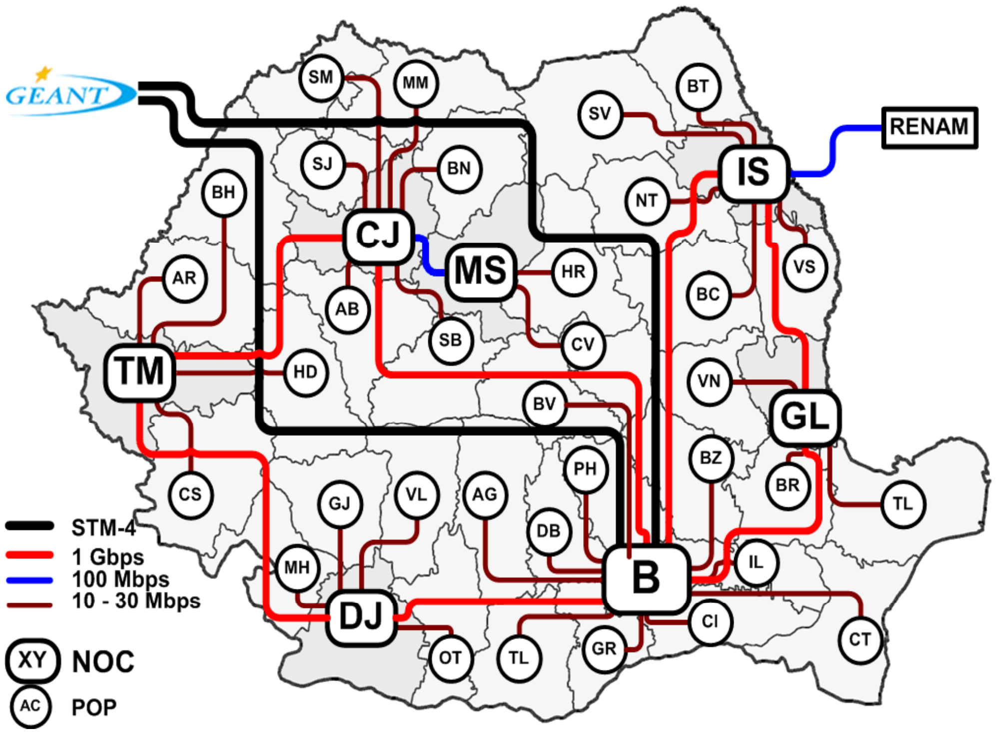

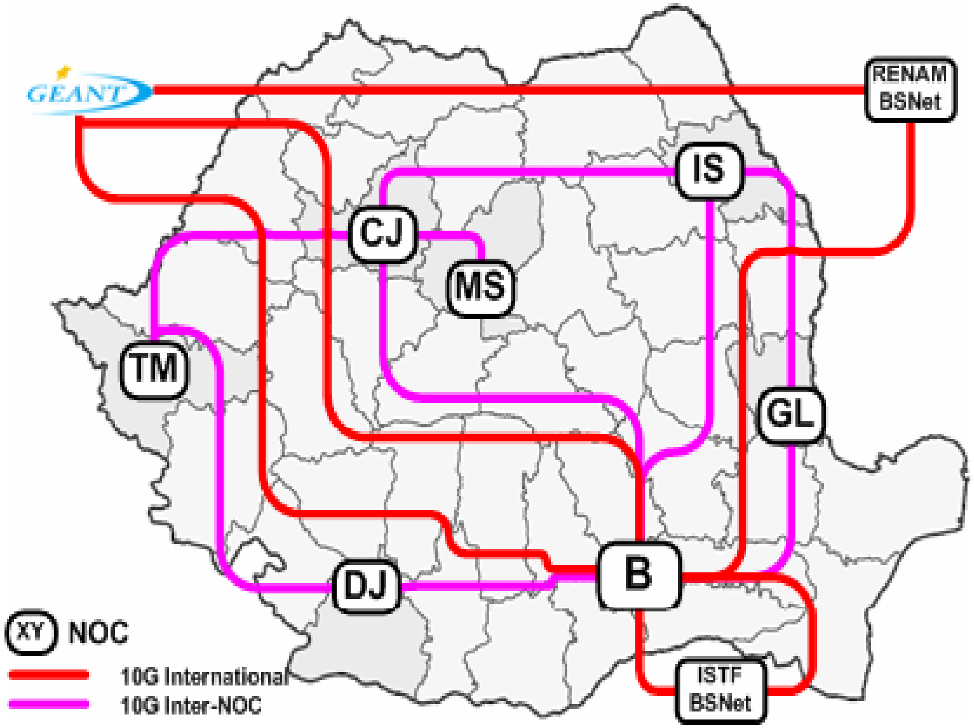

5 Years Since 2002 SCIC Grid Workshop and WSIS Pan-European Ministerial Meeting in Bucharest

International ICFA Workshop on Digital Divide Issues

UNIVERSIDAD IBEROAMERICANA

http://fismat.uia.mx/HEP/ICFADDW2007

ICFA SCIC

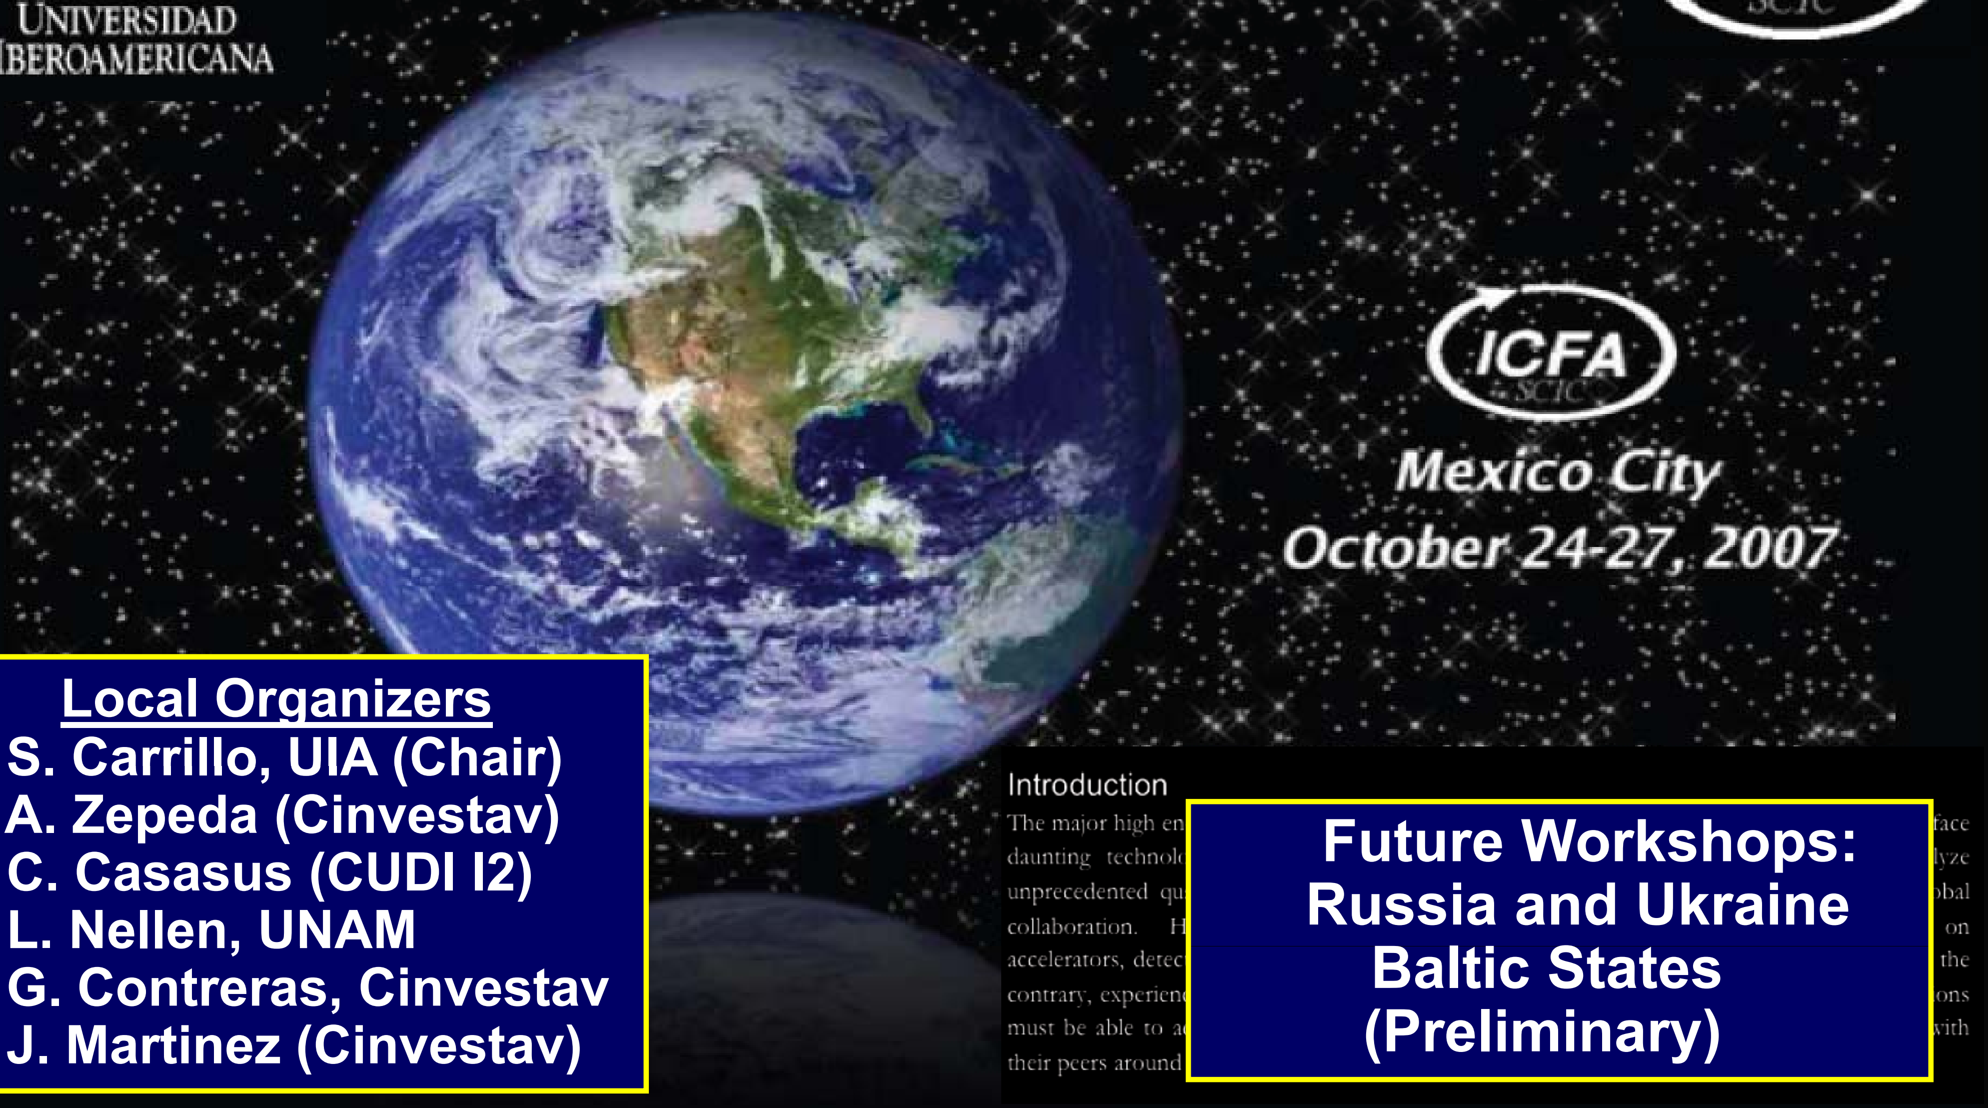

Local Organizers

S. Carrillo, UIA (Chair)
A. Zepeda (Cinvestav)
C. Casasus (CUDI I2)
L. Nellen, UNAM
G. Contreras, Cinvestav
J. Martinez (Cinvestav)

Future Workshops:
Russia and Ukraine
Baltic States
(Preliminary)

Latin Am. Country	NREN Organization	National connections	External Capacity A.Santoro
Argentina	RETINA	256 Kbps – 34 Mbps	90 Mbps RedCLARA (temp. disconnected 12/2006-1/2007)
Bolivia	ADSIB	64 – 128 Kbps	1.5 Mbps (commodity)
Brazil	RNP	up to 10 Gbps	1 Gbps (incl. RedCLARA)
Chile	REUNA	155 Mbps	90 Mbps RedCLARA
Colombia	RENATA	10 Mbps	10 Mbps RedCLARA
Costa Rica	CR2Net	32 – 512 Kbps	10 Mbps RedCLARA (currently disconnected)
Cuba	RedUniv	19.2 Kbps– 2Mbps	Not known (commodity)
Ecuador	CEDIA		10 Mbps RedCLARA
El Salvador	RAICES		10 Mbps RedCLARA
Guatemala	RAGIE		10 Mbps RedCLARA
Honduras	HONDUnet		Not known (commodity)
Mexico	CUDI	155 Mbps	1 Gbps-RedCLARA, 1 Gbps CENIC (PacWave)
Nicaragua	RENIA	100 Mbps	10 Mbps RedCLARA
Panama	PANNET/ SENACYT	2 Mbps	45 Mbps RedCLARA
Peru	CONCYTEC	2 Mbps	45 Mbps RedCLARA
Uruguay	RAU	64 Kbps to 1 Mbps	34 Mbps RedCLARA
Venezuela	REACCIUN	155 Mbps & 34 Mbps	90 Mbps RedCLARA and Ampath

Summary and Conclusions

- Our major networks backbones and int'l links are transitioning from 10G to N X 10G; Progressing much faster than Moore's law
- HEP Network Requirements Roadmaps: ~300X-1000X/decade
- We continue to learn to use long N X 10G network paths effectively
- Effective use of the Global ensemble of T0/T1/T2 centers requires
  - Managing the network as a resource, along with CPU & storage
  - Developing a new generation of dynamic circuit services: GLIF; UltraLight/VINCI, with λStation, TeraPaths, OSCARs
  - New technologies: VCAT/LCAS, a global system-level view
- We will have to coordinate, with a coherent vision, and corresponding development, deployment, and O&M plans
- As we progress it is essential to address the problem of the Digital Divide in our community
  - To form a truly global partnership for leading-edge science
  - An ongoing commitment, and ongoing work is required