



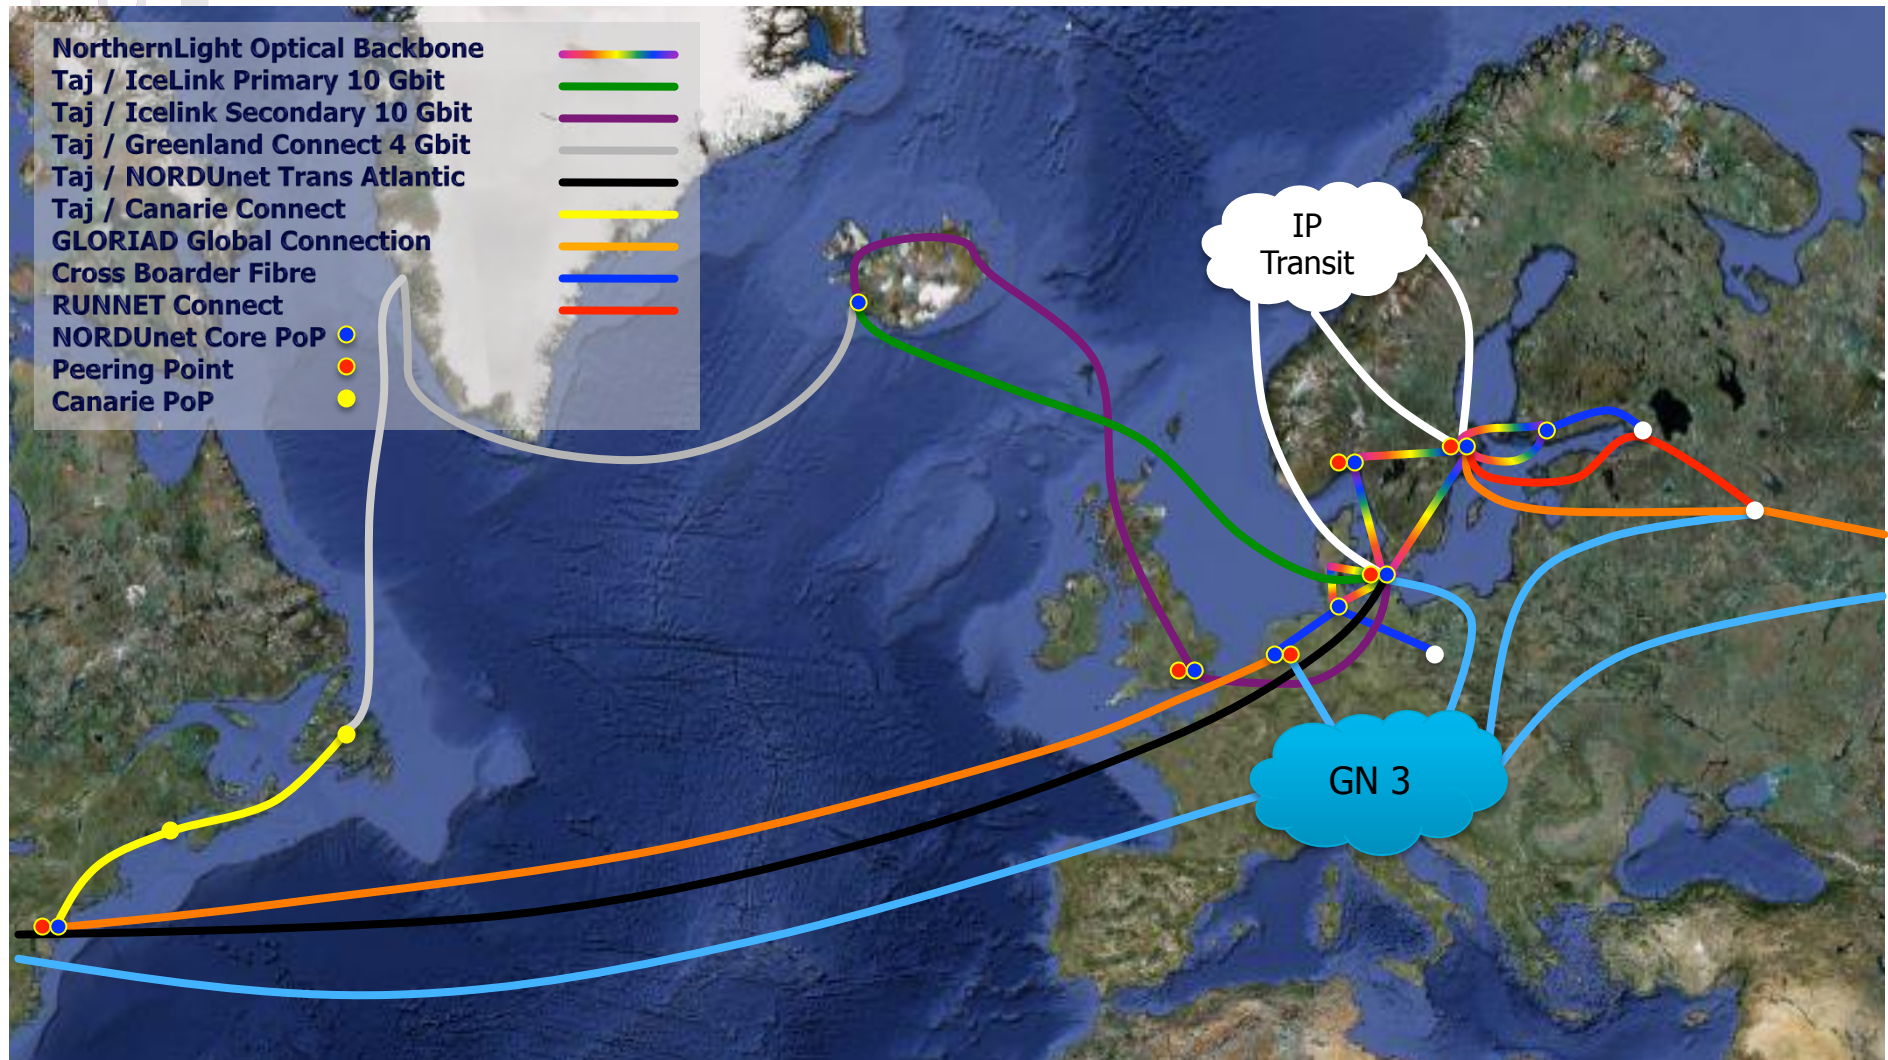
Integration of GLIF services into the Grid middleware

Lars Fischer

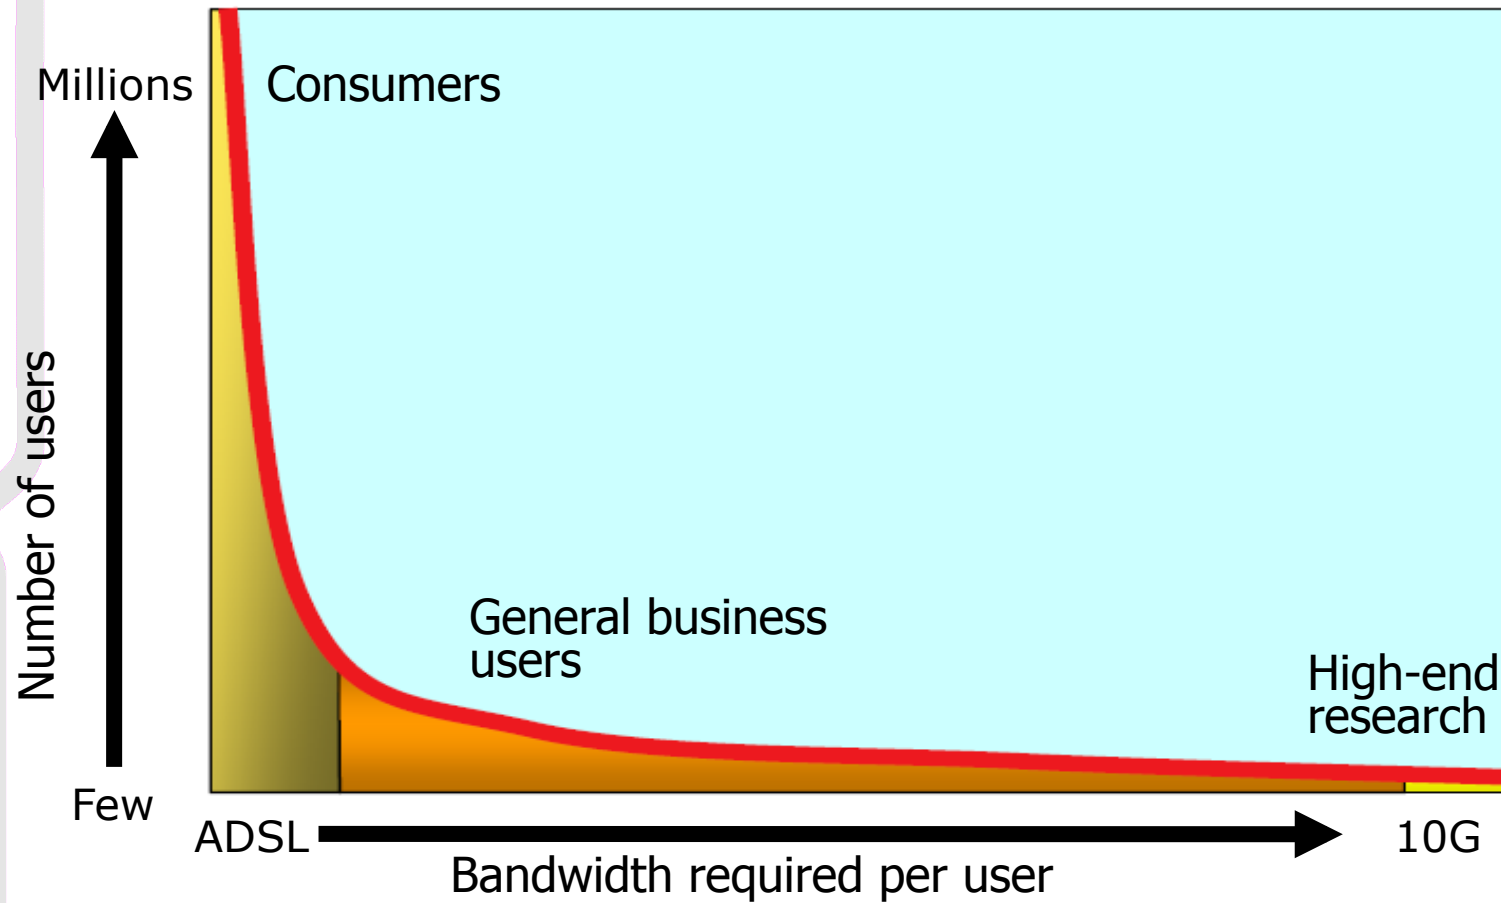
TERENA NRENs-n-GRID workshop

Barcelona, 22 September 2009

- Regional Network for Nordic countries
 - Denmark, Sweden, Iceland, Norway, Finland
- Services
 - Network Interconnect,
 - International Connectivity
 - Coordination
 - Representation, international projects and relations
- Organization
 - Jointly owned by Nordic NRENs
 - Funded according to GNP







- Offload heavy users onto (simple) private circuit-switched networks
 - Save capacity and cost on routed networks
 - As a bonus, private networks provide predictable network properties (i.e., true QoS)
- Challenges
 - Coordination of provisioning, routing
 - Multi-domain lambdas
 - International connectivity
 - Capability to create lambdas on the fly

NORDUnet

Nordic infrastructure for Research & Education

2003 Lambda Grid Workshop



- Global Lambda Integrated Facility

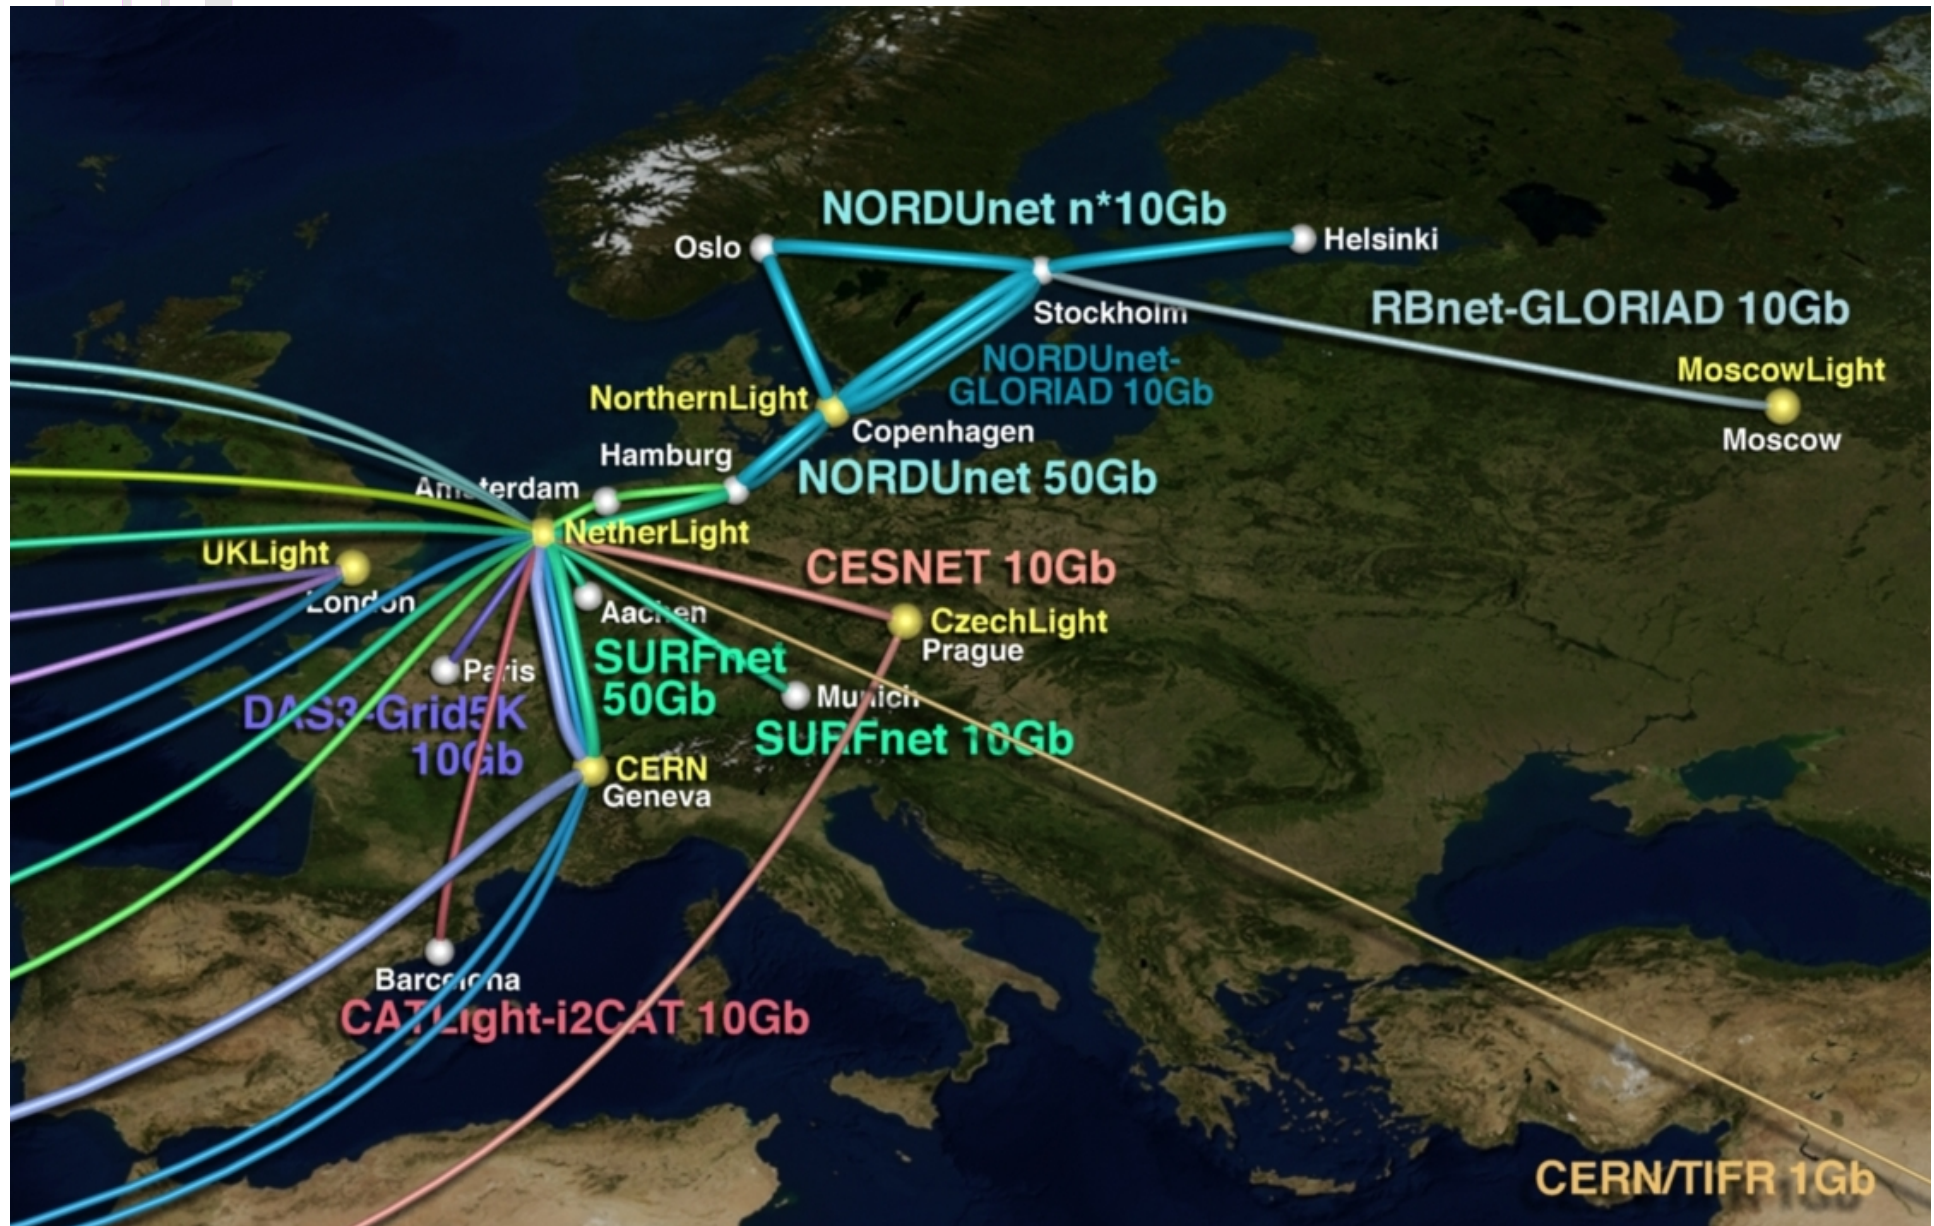


- GLIF is an international, virtual organization that promotes the paradigm of lambda networking
- The activities of GLIF are two-fold:
 - GLIF participants jointly make lambdas available as an integrated global facility for use in data-intensive research
 - GLIF brings together leading networking engineers worldwide, who collaborate to bring the technology forward and pioneer new applications – a world-scale laboratory

- GLIF is open
 - to any organization that shares the GLIF vision and is willing and able to make resources (e.g., equipment and lambdas) available on an agreed basis when they are not required for its own needs.
 - GLIF is open to organizations whose experts contribute actively to the technical work in the GLIF working groups.
- Lightweight governance structure.
 - Managed as cooperative activity with partners rather than members
 - Short-lived taskforces
- Secretariat functions provided by TERENA with contributions from sponsors.

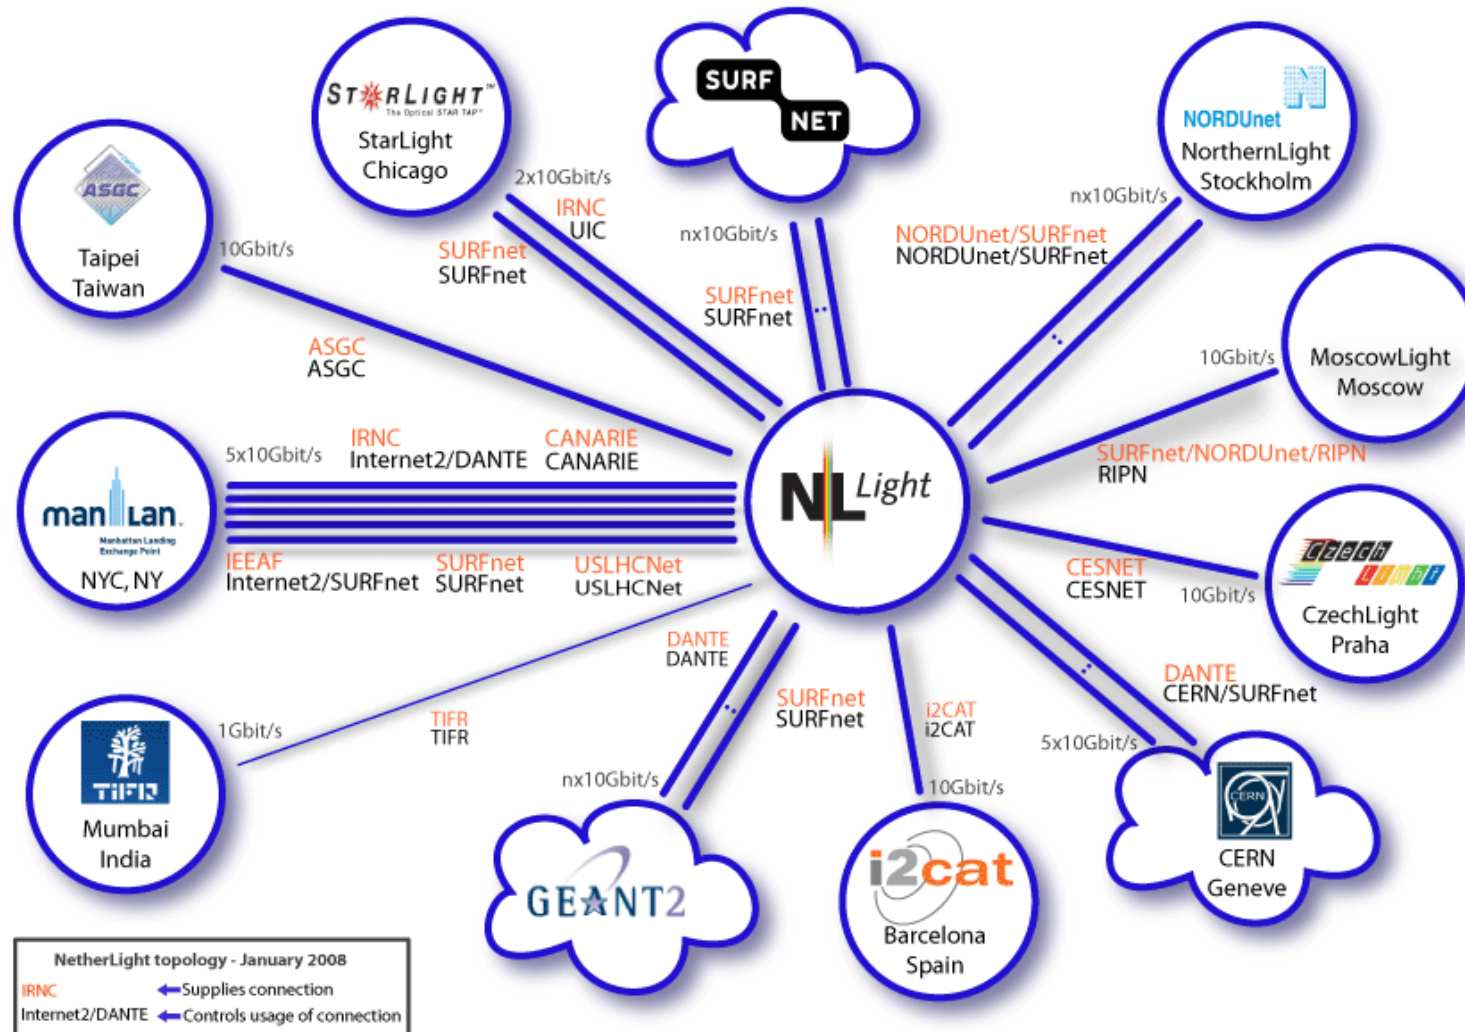
- **AARNet**
- **AMPATH**
- **ANL/MCS**
- **Calit2**
- **CANARIE**
- **CENIC**
- **CERN**
- **CERNET**
- **CESNET**
- **CPqD**
- **DFN**
- **FAST**
- **Fermilab**
- **GLORIAD**
- **HEAnet**
- **i2CAT**
- **iCAIR**
- **IEEAF**
- **Indiana Univ.**
- **Internet2**
- **JGN-II**
- **JANET(UK)**
- **KRLight**
- **MIT**
- **MCNC**
- **MAX**
- **NLR**
- **Northwestern**
- **NSF**
- **NetherLight**
- **NORDUnet**
- **Nortel**
- **PNWGP**
- **RNP**
- **SARA**
- **StarLight**
- **SURA**
- **TeraGrid**
- **TERENA**
- **TransLight**
- **TWAREN**
- **Univ. Amsterdam**
- **UCL**
- **UIC**
- **WIDE**
- **Univ. Maryland**
- **Univ. Washington**
- **UltraLight**

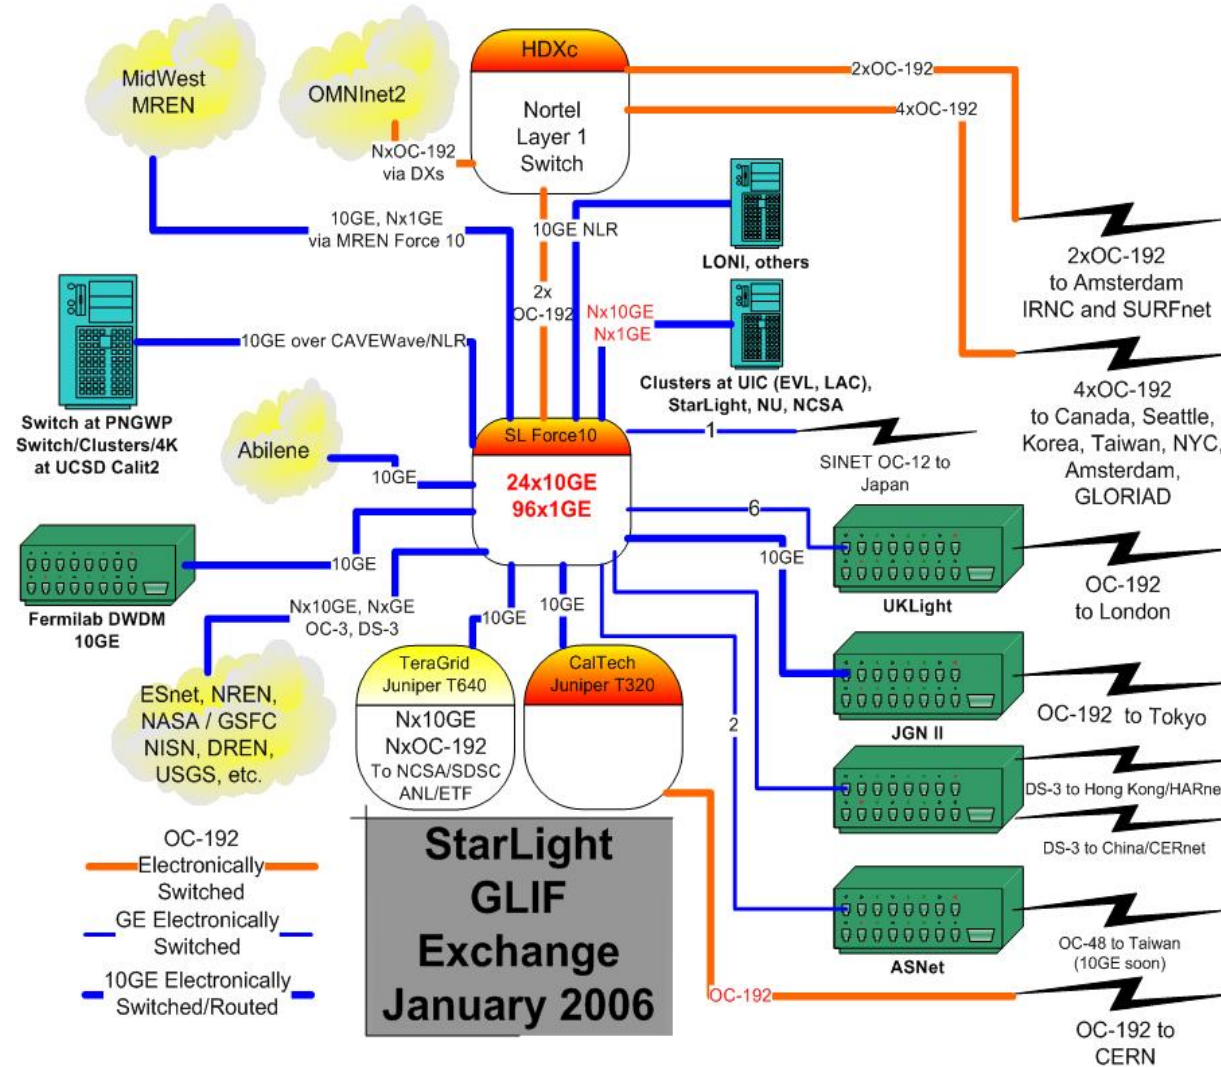




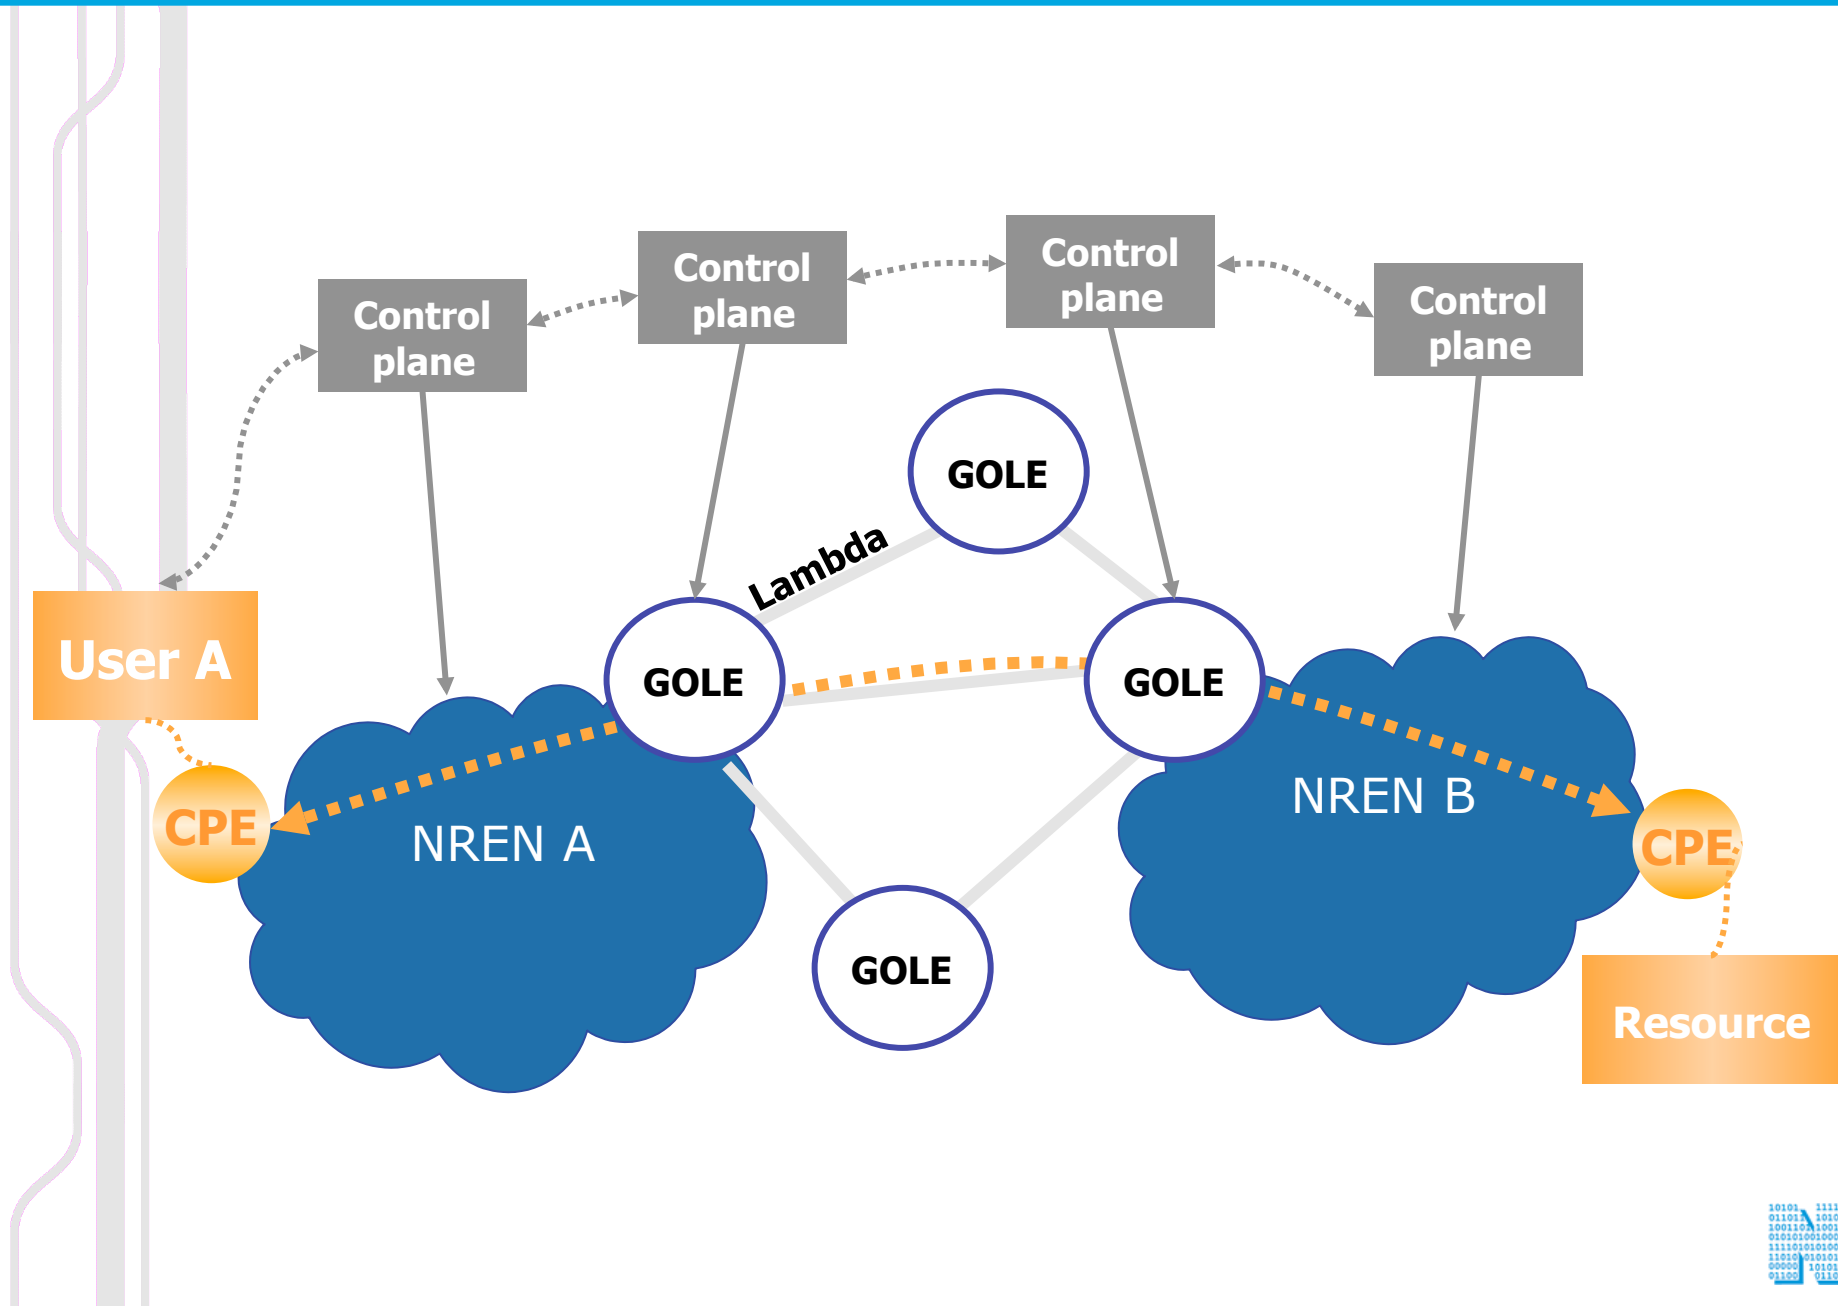
- **GOLE**: GLIF open lightpath exchange, comprised of one or more network devices performing lightpath switching.
- **lambda**: a high-capacity (optical) circuit which terminates on a GOLE
- **lightpath** [service]: a high capacity circuit or QoS-supported virtual circuit (or the concatenation of several of these to form an end-end lightpath.) A lightpath can take up a portion or all of the capacity of a lambda.
- **GLIF network resources**: GOLE's and lambdas

- The key infrastructure of the GLIF collaboration
- Open Exchange Points for circuit-switched networks – in the tradition of IP Internet Exchanges
- Everyone can bring a circuit and exchange traffic with everyone else
- A GOLE is typically a SDH/SONET switch matrix, possibly with Ethernet and photonic switching
- Each GOLE will have lambdas connecting it to other GOLEs



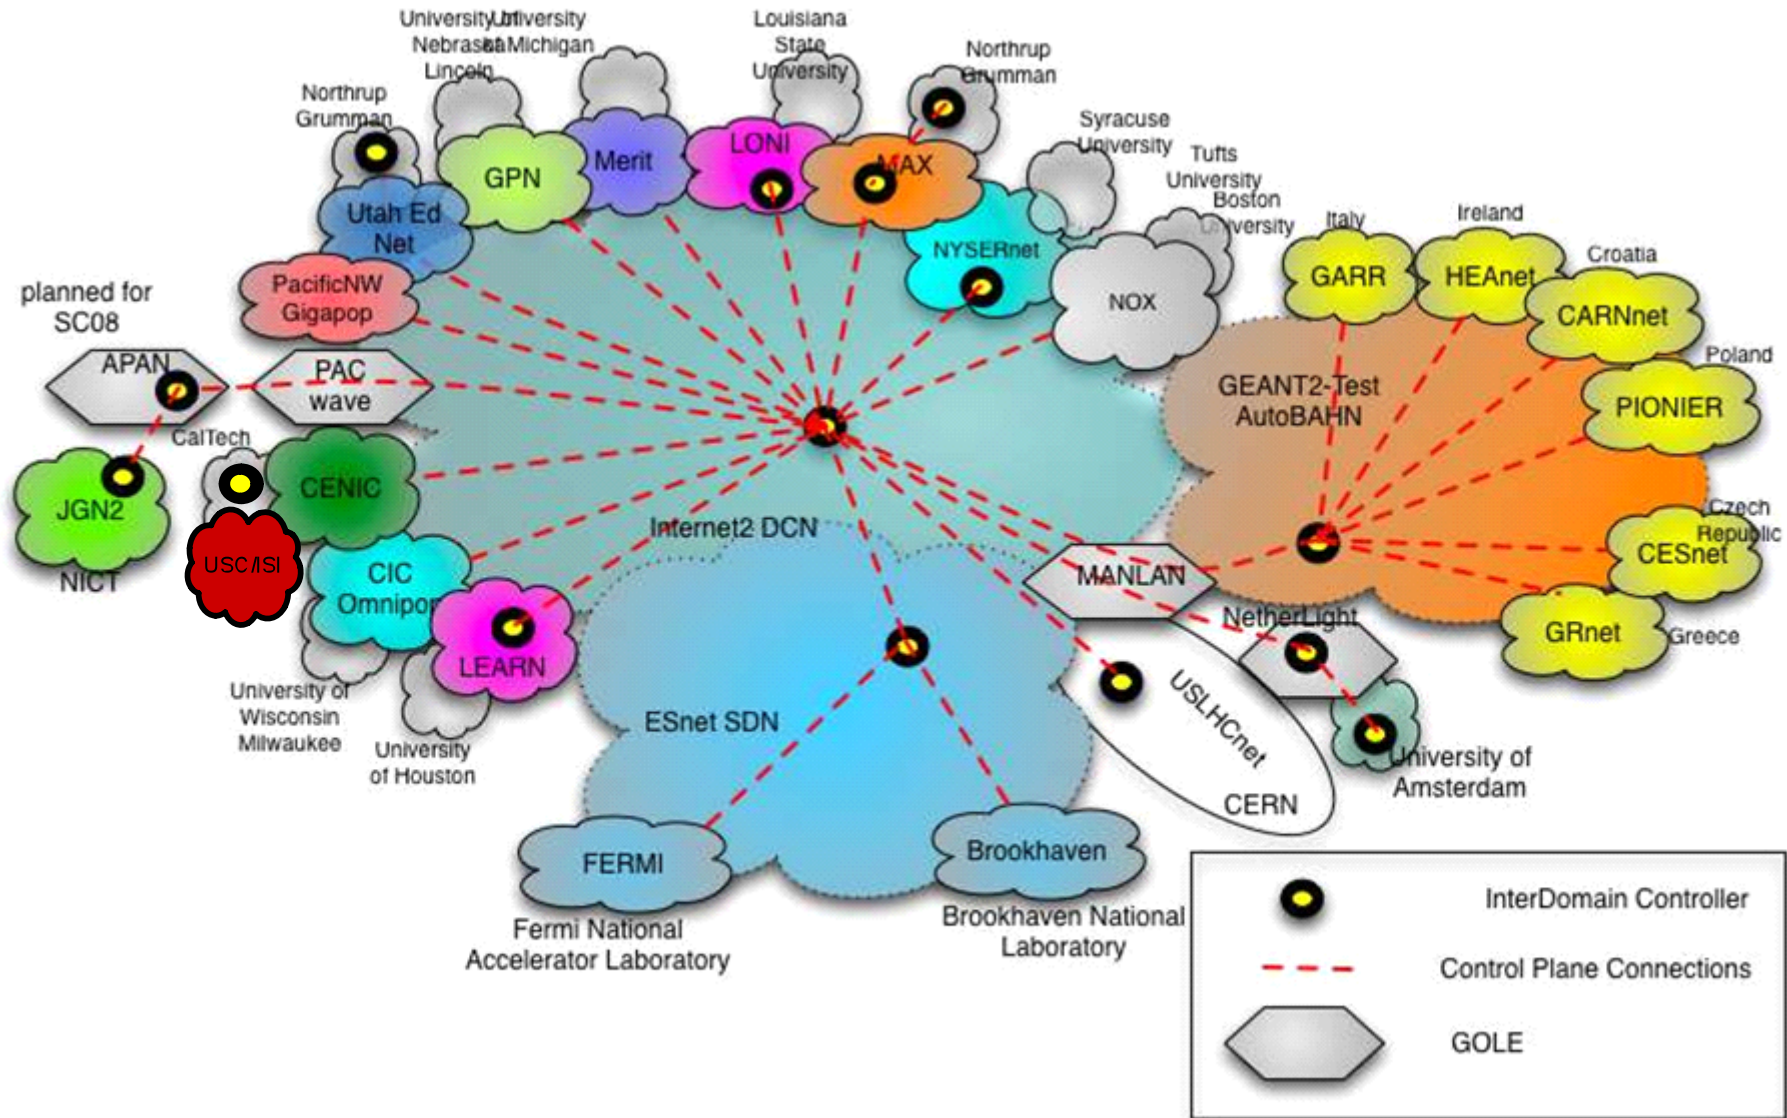


- In GLIF, end-to-end connectivity is constructed by piecing together circuits from GOLE to GOLE, eventually connecting end-sites
- There are GOLEs all over the globe
- GOLEs are typically operated by NRENs
- GOLE operators have coordination meeting
- Fiber and circuits are deployed to connect GOLEs



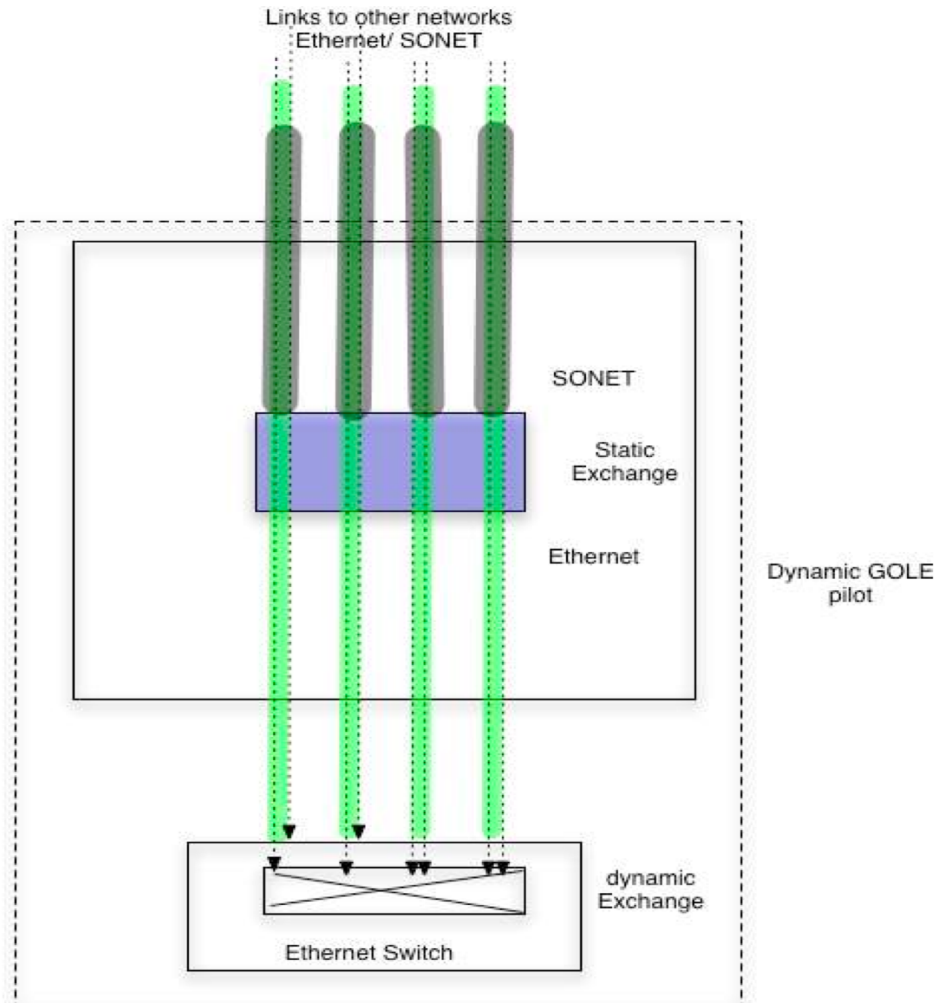
- Dynamic Lightpaths
 - Some lightpaths are short-lived
 - Some applications do late resource binding
 - Allow allocation, setup, and tear-down of circuits through automated system
 - Static allocation wasteful of resources
- Dynamic control plane systems
 - DRAC
 - Harmony
 - Oscars
 - DRAGON
 - AutoBAHN
 - ...



- Users need ubiquitous end to end dynamic lightpath connectivity over a multi-domain infrastructure
- How do we glue things together?
 - Centralistic models will be too complex and will not scale
 - Cooperation between networks requires a interworking control plane
- Create a loose cooperation between domains
 - Standardized interfaces between domains
 - Create user friendly AAA features
- Paving the way to a ubiquitous and scalable Services Grid



- Give users the ability to modify cross connects at the GOLE without intervention by GOLE staff
 - Cross-connecting dynamic lightpaths
 - Interconnection point for dynamic lighthpath service domains
- Adaptations between cross connected segments
- GLIF Dynamic GOLE WG
 - Defining dynamic GOLE capabilities is an ongoing project
 - Ongoing work also in OGF NSI WG

- Pilot of Dynamic GOLE capabilities
 - Focus on Ethernet VLAN cross-connects
 - Deploy web services interface for user creation of cross connects
 - Support IDC interoperability cross-connect control plane
 - Attached to existing full-service GOLE
- Test-bed for experimental services
 - Interconnect networks to create dynamic circuits, using web services interface between network and GOLE control plane
 - Engage experimental applications
 - Connect experimental application sites

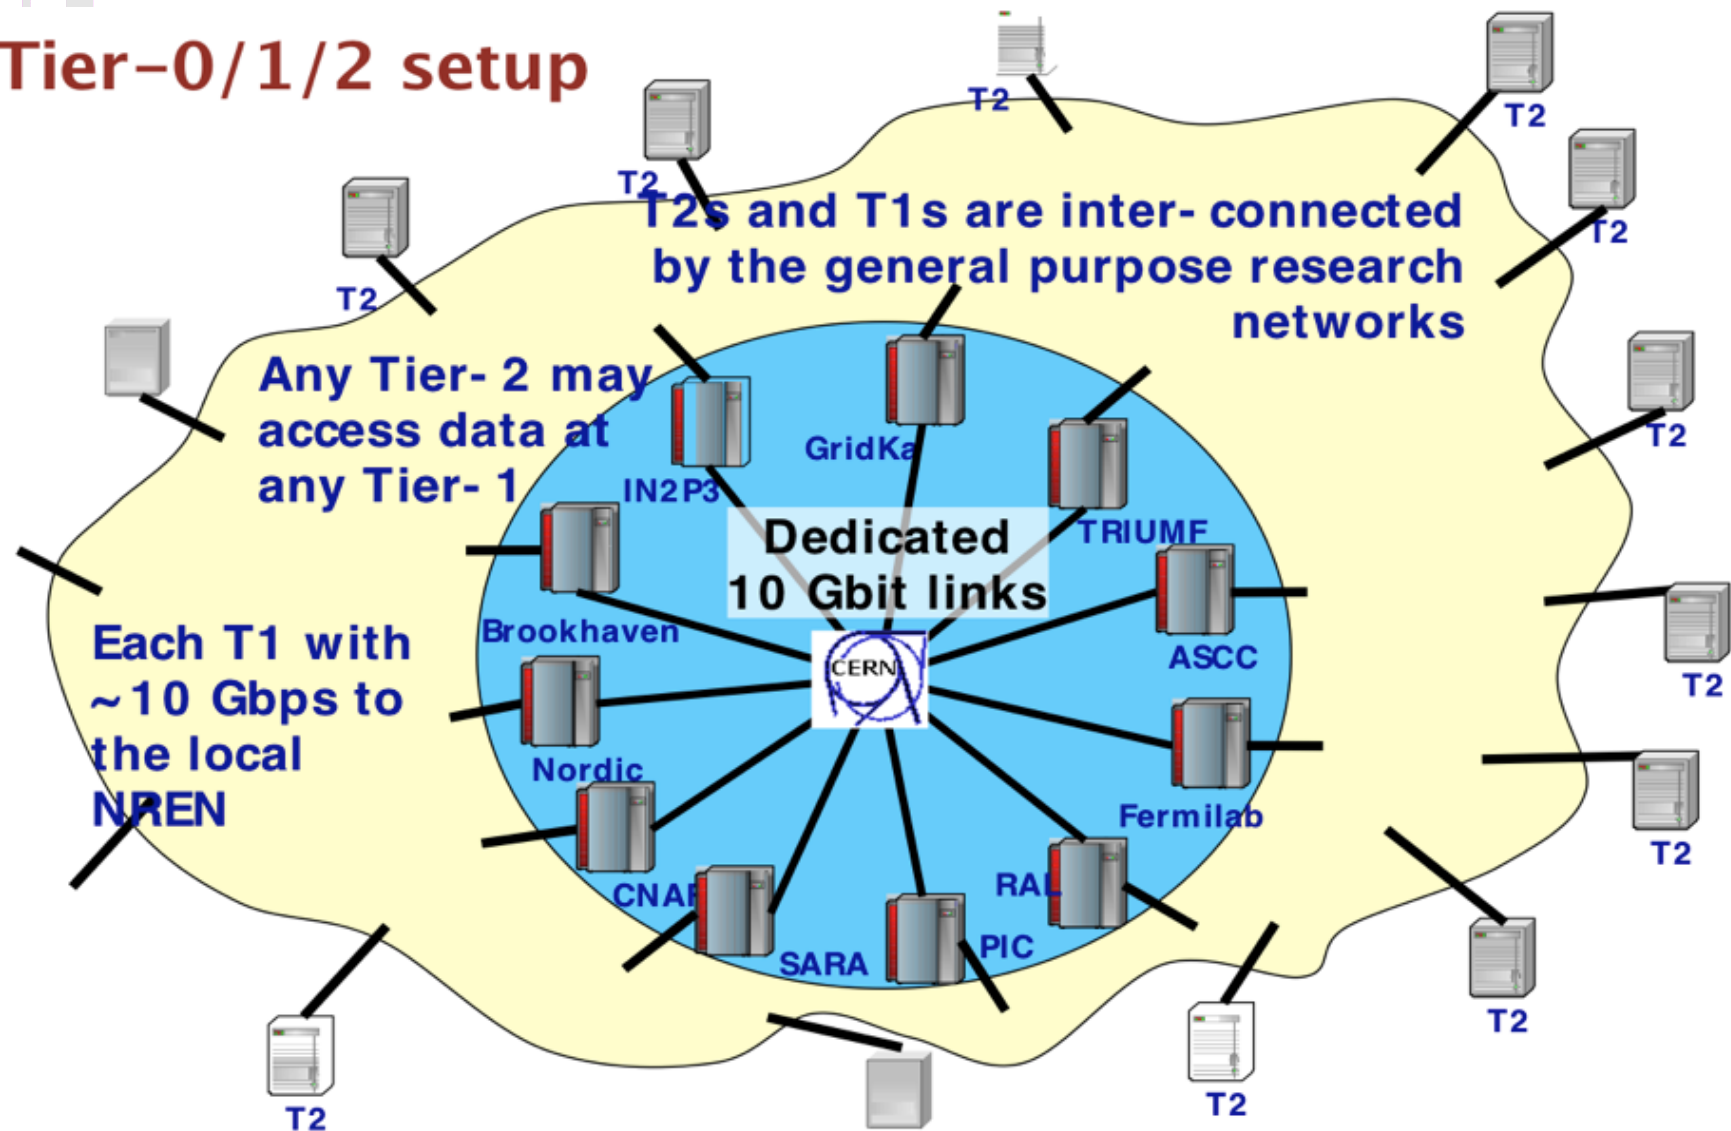


-  SONET Connection to remote Network
-  Ethernet Link using a subset of SONET
-  VLAN carried in Ethernet Link



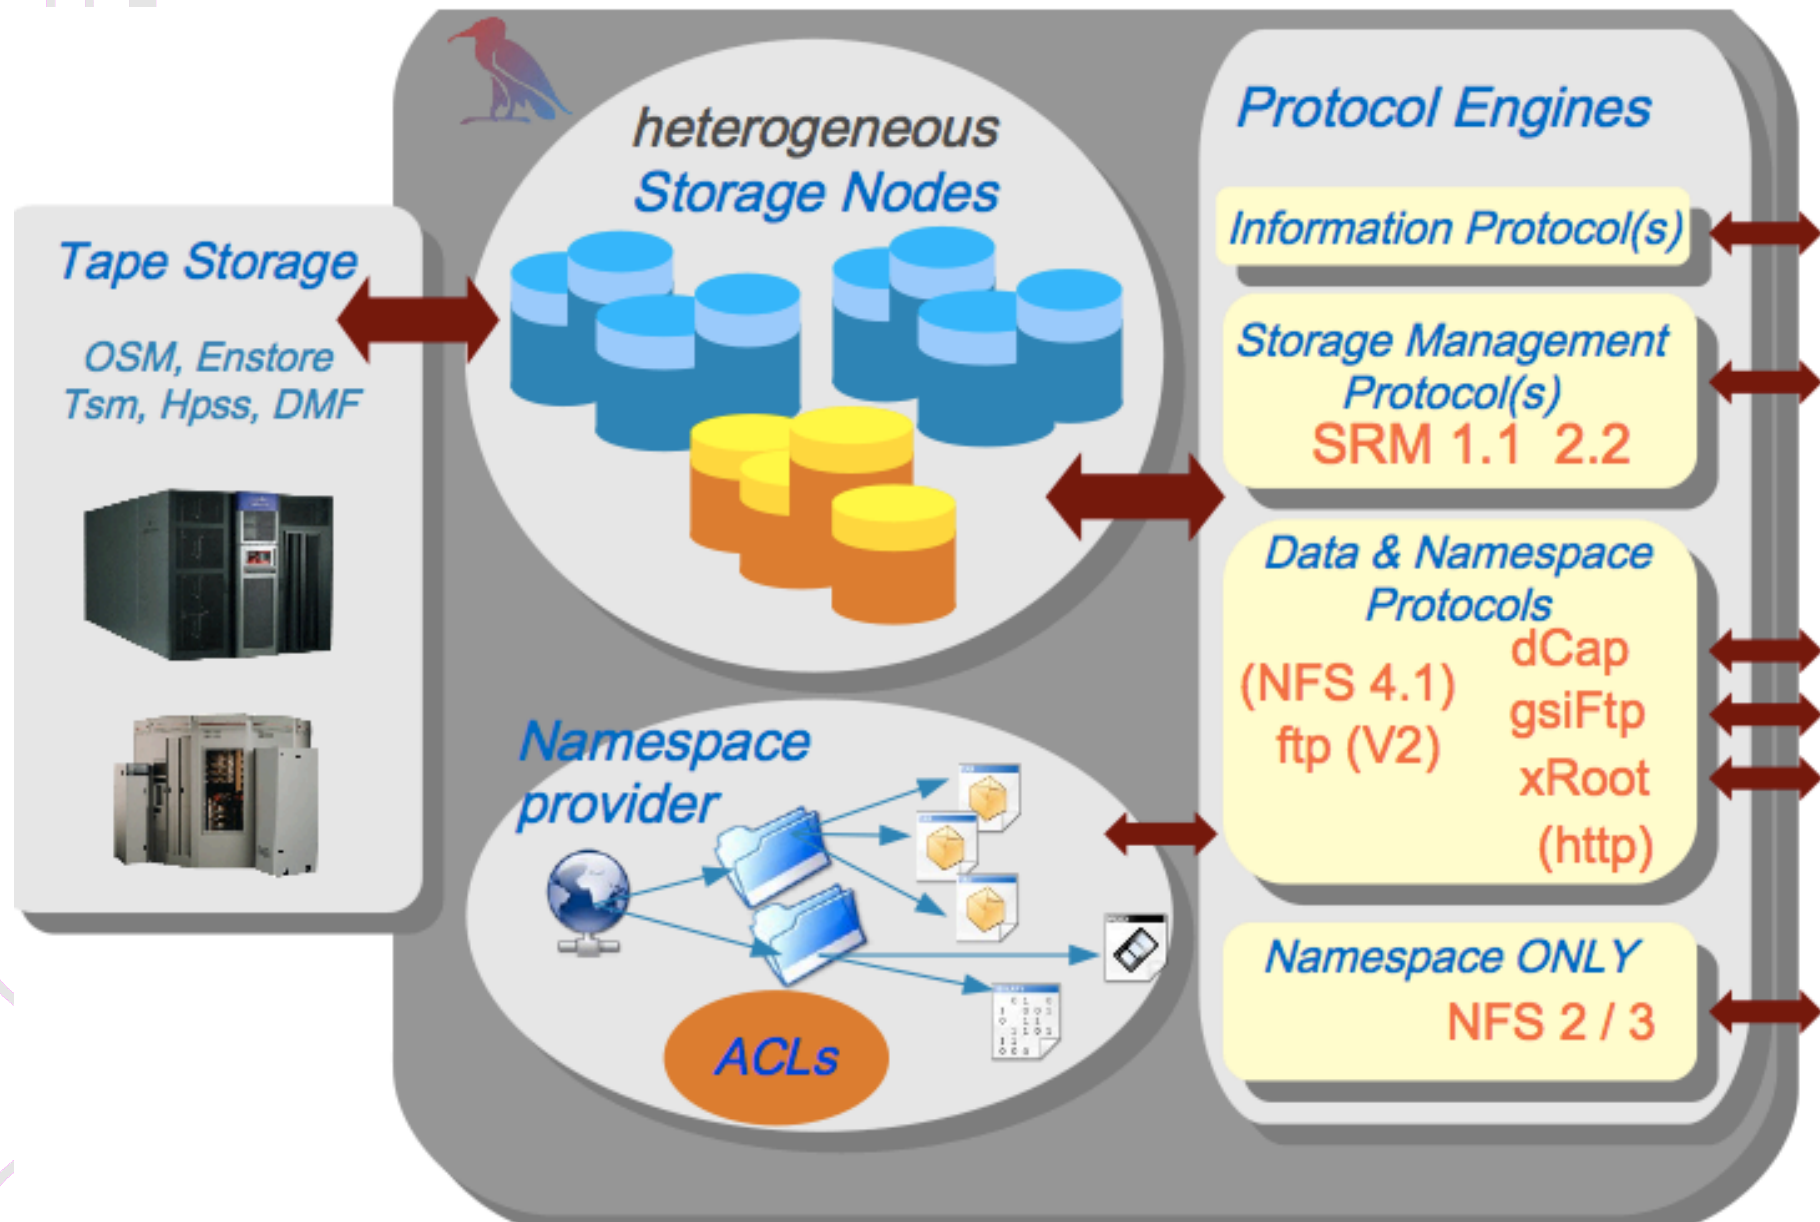


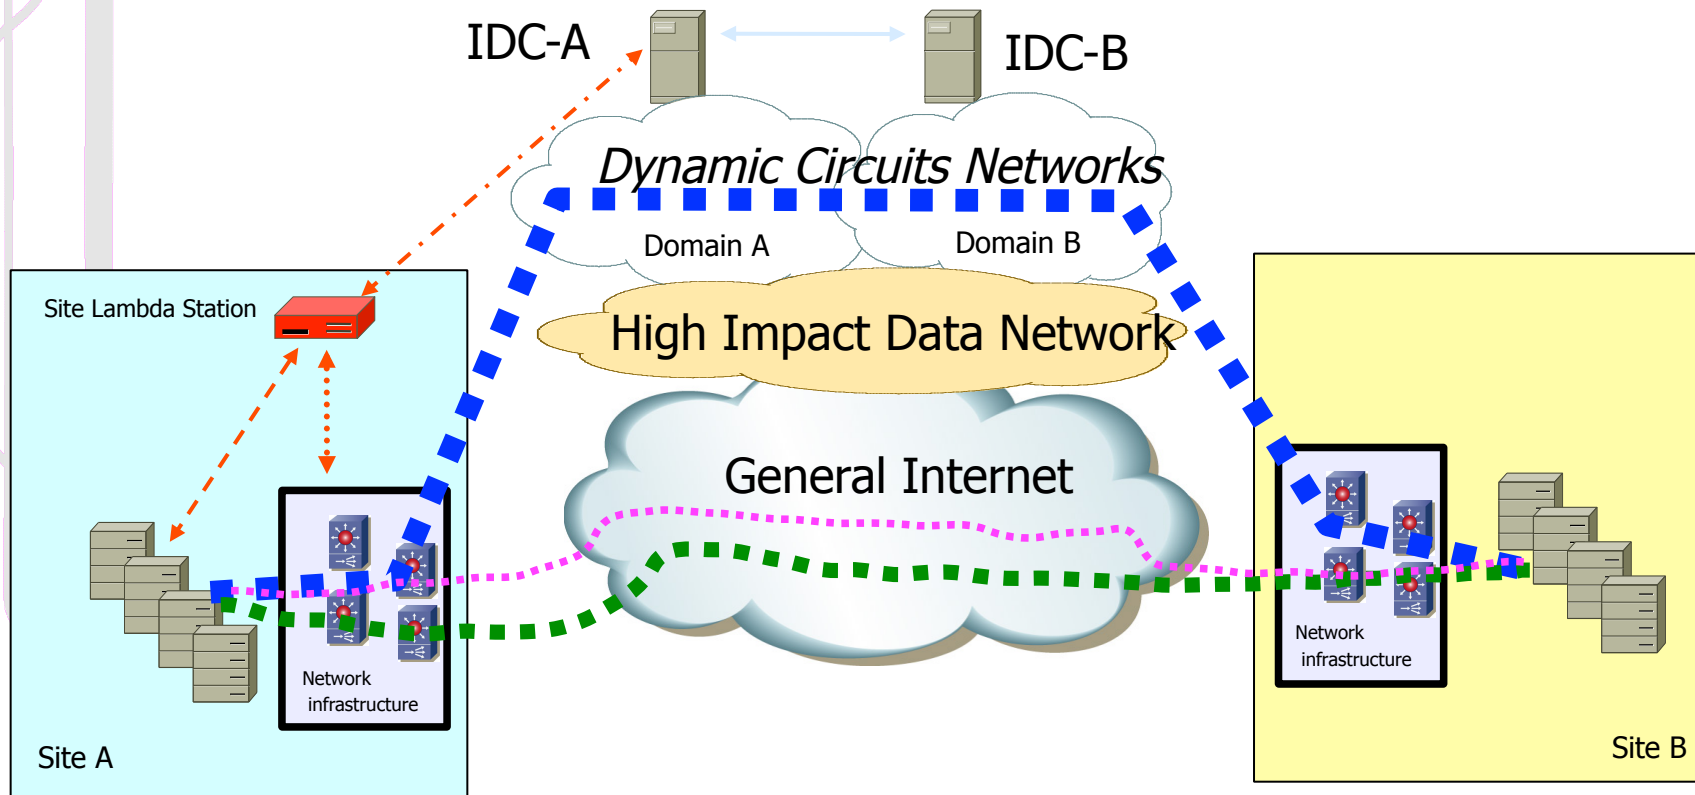
Tier-0/1/2 setup



- LHC OPN is built for on-going flows
 - T0 – T1 data transfers
 - T1 – T1 transfers (some better connected than others)
- Does not connect Tier-2 sites
 - Tier-2's (in general) connected over public IP
 - Some Tier-2 have private connection to local Tier-1
 - Can restrict capacity for Tier-2 to remote Tier-1 traffic
- T0/T1/T2 workgroup looking into design for future needs
 - Connecting to GOLE's may be significant

Tier-1 Storage: dCache





- ■ ■ ■ ■ ■ ■ ■ Default network path
- ■ ■ ■ ■ ■ ■ ■ Forward path across DCN
- - - - - Reverse path across general Internet

Control plane

- - - - - Circuit call setup & teardown
- - - - - LAN reconfiguration to use circuit
- - - - - Application-initiated request/response



- Facilities
 - Lambda Station can talk to Oscars / DCN
 - dCache can talk to Lambda Station
 - DCN can exploit Pilot dynamic GOLE
- Remote Tier1 - Tier2 transfer
 - Initiate transfer with dCache middleware
 - Use Lambda Station to redirect flow and use DCN to provision dynamic lightpath
 - Create connectivity across connected Dynamic GOLEs
- Has been demonstrated for US-based Tier2 transfers w/Oscars (single domain)



- Many alternatives to be explored
 - Initiate dynamic lightpath directly from dCache
 - Initiate dynamic lightpath from grid middleware, above dCache layer
 - Initiate dynamic lightpath from application layer
- Towards Network Services Interface
 - We need prototypes, application experience
 - Different setups for different applications
 - Extend beyond VLANs

- 9th Annual Global Lambda Workshop
- 27-28 October 2009, Daejeon, Korea
- Daejeon Convention Center
- Hosted by KISTI



Questions?

lars@nordu.net

9/20/09

