The DRAGON Project and Application Controlled Networks

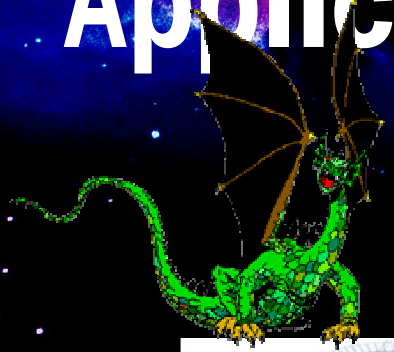

A future space of hybrid application oriented service networks

Jerry Sobieski · Mid-Atlantic Crossroads
Tom Lehman University of Southern California/ ISI-East
Bijan Jabbari George Mason University
Don Riley University of Maryland

Presented at
IPOP 2006
Tokyo, Japan
June 23, 2006

The Single-Slide Overview

- DRAGON
  Dynamic Resource Allocation over GMPLS Optical Networks
- US National Science Foundation four year $7M “Experiemental Infostructure Networks” (EIN) program
  - Extends from Sept 03 thru Aug 07
  - Testbed deployed in the Washington DC region…
- Objectives:
  - Develop GMPLS based dynamic LightPath signaling and routing capabilities for the R&E community
  - Employ all-photonic [dynamic] wave services within a metro scale network
  - Integrate these capabilities with real applications in order to promote both the capabilities of the application and the efficacy of these network services

Participating Organizations

- **Mid-Atlantic Crossroads (MAX)**
- **USC/ Information Sciences Institute (ISI-East)**
- **George Mason University (GMU)**
- **University of Maryland (UMCP)**
- MIT Haystack Observatory
- NASA Goddard Space Flight Center (GSFC)
- Movaz Networks (commercial partner)
- NCSA ACCESS
- US Naval Observatory
- **Internet2 HOPI**
- **NASA Ames Research Center (AMES)**
- **Univ of Maryland Baltimore County (UMBC)**
- **Northrop Grumman Corp.**
- **Others in the works…**

Project Features and Objectives

- **DRAGON**
  Dynamic Resource Allocation over GMPLS Optical Networks
- US National Science Foundation four year $7M “Experiemental Infostructure Networks” (EIN) program
  - Extends from Sept 03 thru Aug 07
  - Testbed deployed in the Washington DC region…

- Utilize GMPLS protocols for dynamic provisioning of Light Paths
  - Addition of CSPF Path Computation algorithms for wavelength routing
- Inter-domain service routing techniques
  - Network Aware Resource Broker (NARB) for service advertising, inter-domain ERO generation, AAA
- All-Optical metro area network
  - Eliminate OEO in the core, allow alien waves in.
- Application Specific Topology Description Language
  - Formalized means to describe the application topology and network service requirements
- Integration with real applications:
  - E-VLBI
  - HD-CVAN

DRAGON 3yr baby picture

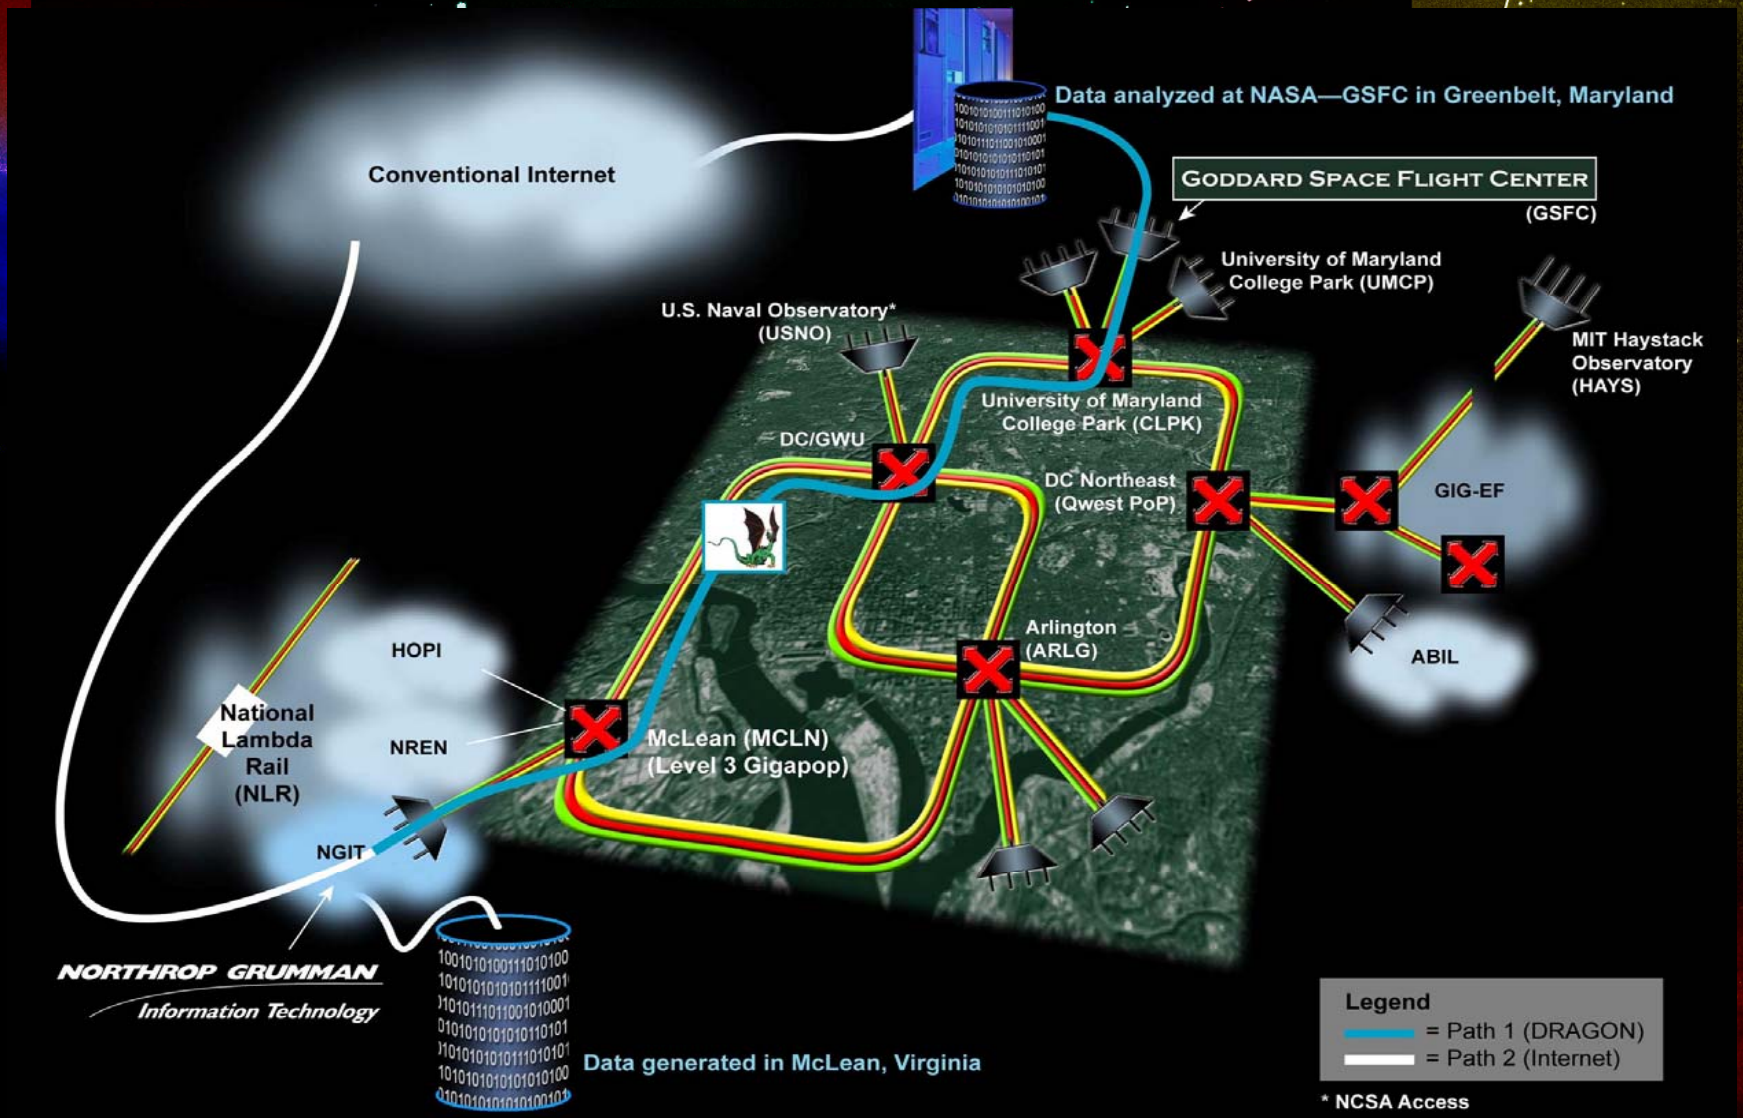

DRAGON Technologies

- Open Source Standards based GMPLS Control Plane:
- Virtual Label Swapping Router – VLSR
  - Open source OSPF-TE & RSVP-TE to control Ethernet switches and fiber switches
- Network Aware Resource Broker – NARB
  - GMPLS-OSPF-TE listener
  - Platform for exchange of inter-domain service routing capabilities for PCE
  - Performs the inter-domain routing, AAA, scheduling, PC
- Advanced "Constrained Shortest Path First" Path Computation Element – CSPF PCE
  - Domain selectable abstraction levels and end-to-end LSP
- Application Specific Topologies - ASTB
  - Formalization of application's resource requirements – particularly the network resources.
- Photonic metro-scale wavelength services
  - Reduce/eliminate unnecessary OEO,
  - allow user generated ITU waves to transit the metro network (alien waves)
  - Framing and encoding agnostic

Virtual Label Switched Router: VLSR

- Many networks consist of switching components that do not speak GMPLS,
  - e.g. current ethernet switches, fiber switches, etc
- The VLSR implements open source versions of GMPLS-OSPF-TE and GMPLS-RSVP-TE and runs on a Unix based PC/workstation
- The VLSR translates GMPLS protocol events into generic pseudo-commands for the covered switches.
  - The pseudo commands are tailored to each specific vendor/architecture using SNMP, TL1, CLI, XML, or a similar protocol.

VLSR Abstraction

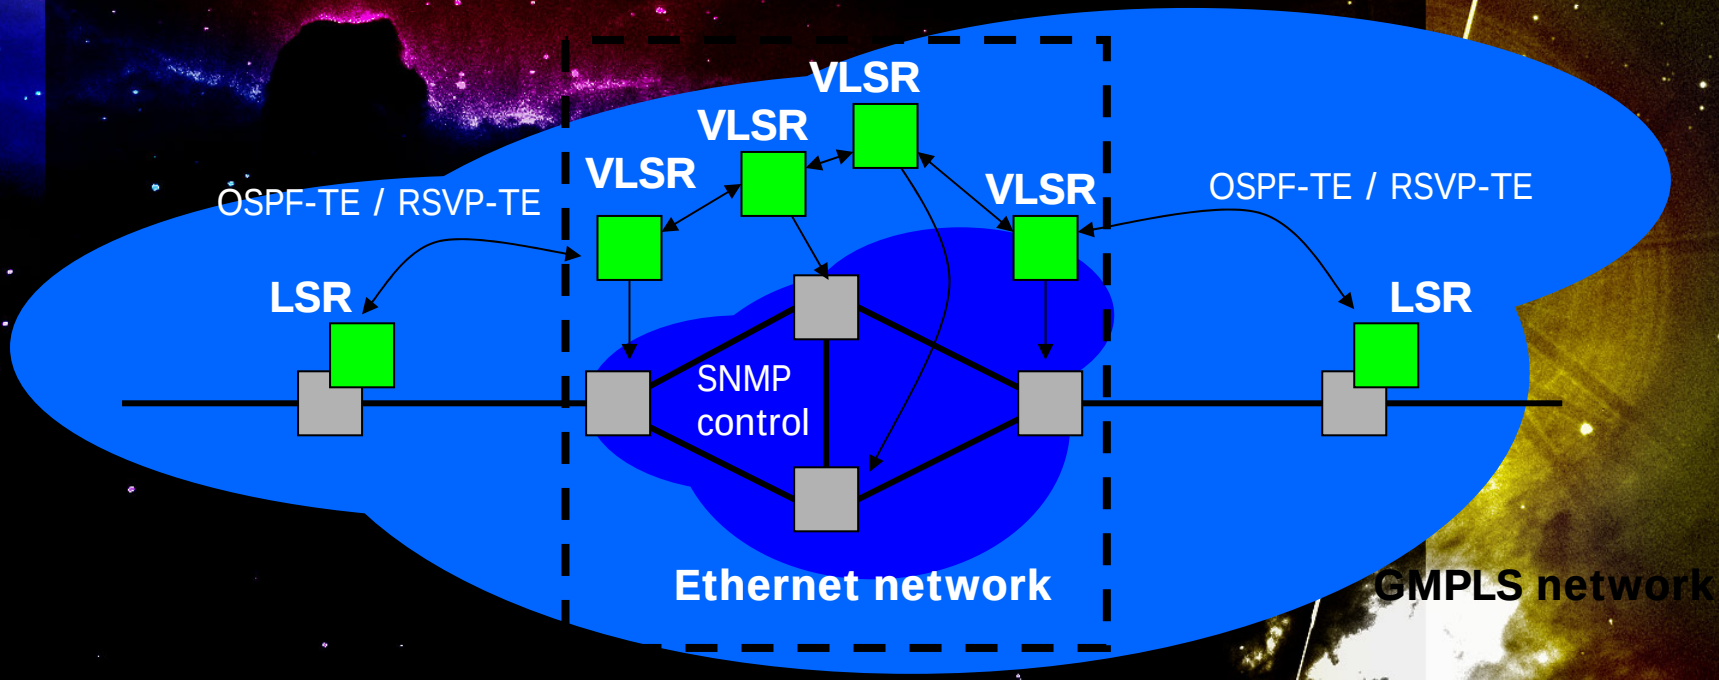

DRAGON Control Plane

Virtual Label Switching Router (VLSR)

- Open source protocols running on PC act as GMPLS signaling entity
- Control PCs participate in protocol exchanges and reprovision covered switch according to protocol events (PATH setup, PATH tear down, state query, etc)

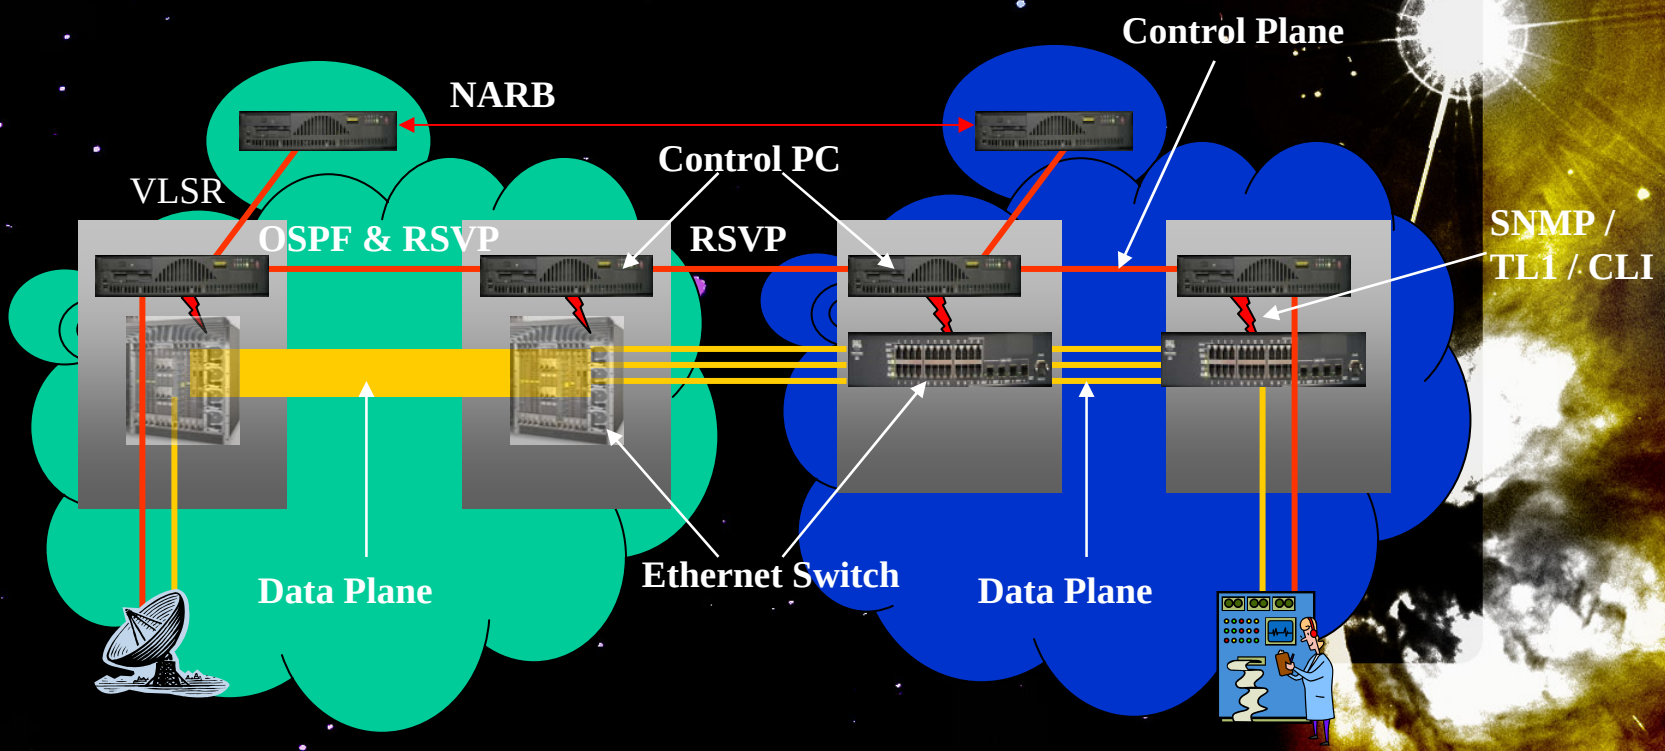

Network Aware Resource Broker: NARB

- Interdomain exchange of service capability information and topology
- Carries a modified TEDB that can support
  - AAA
  - Scheduling
  - PCE and E2E loose hop ERO generation

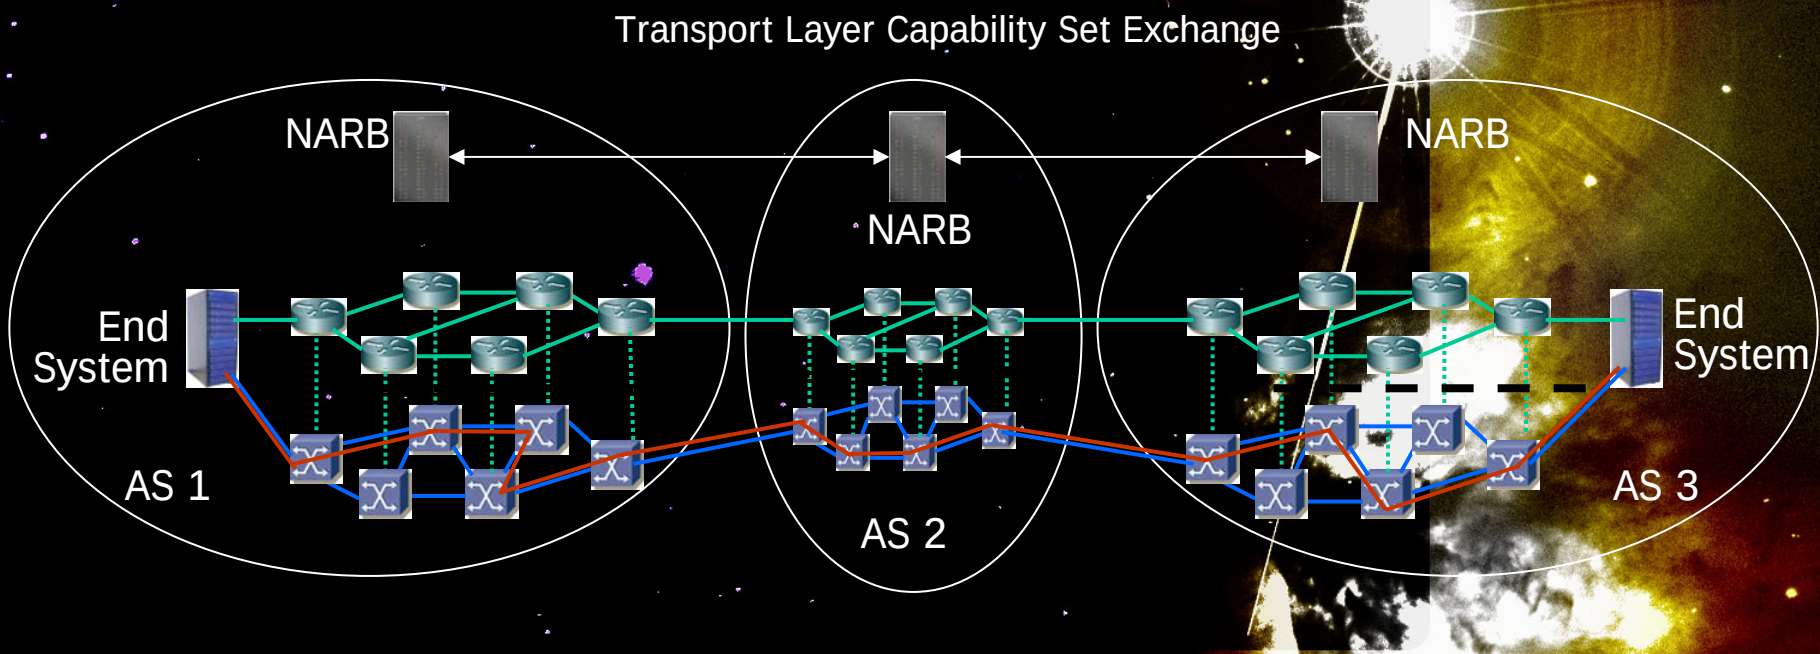

Inter-Domain Topology Summarization

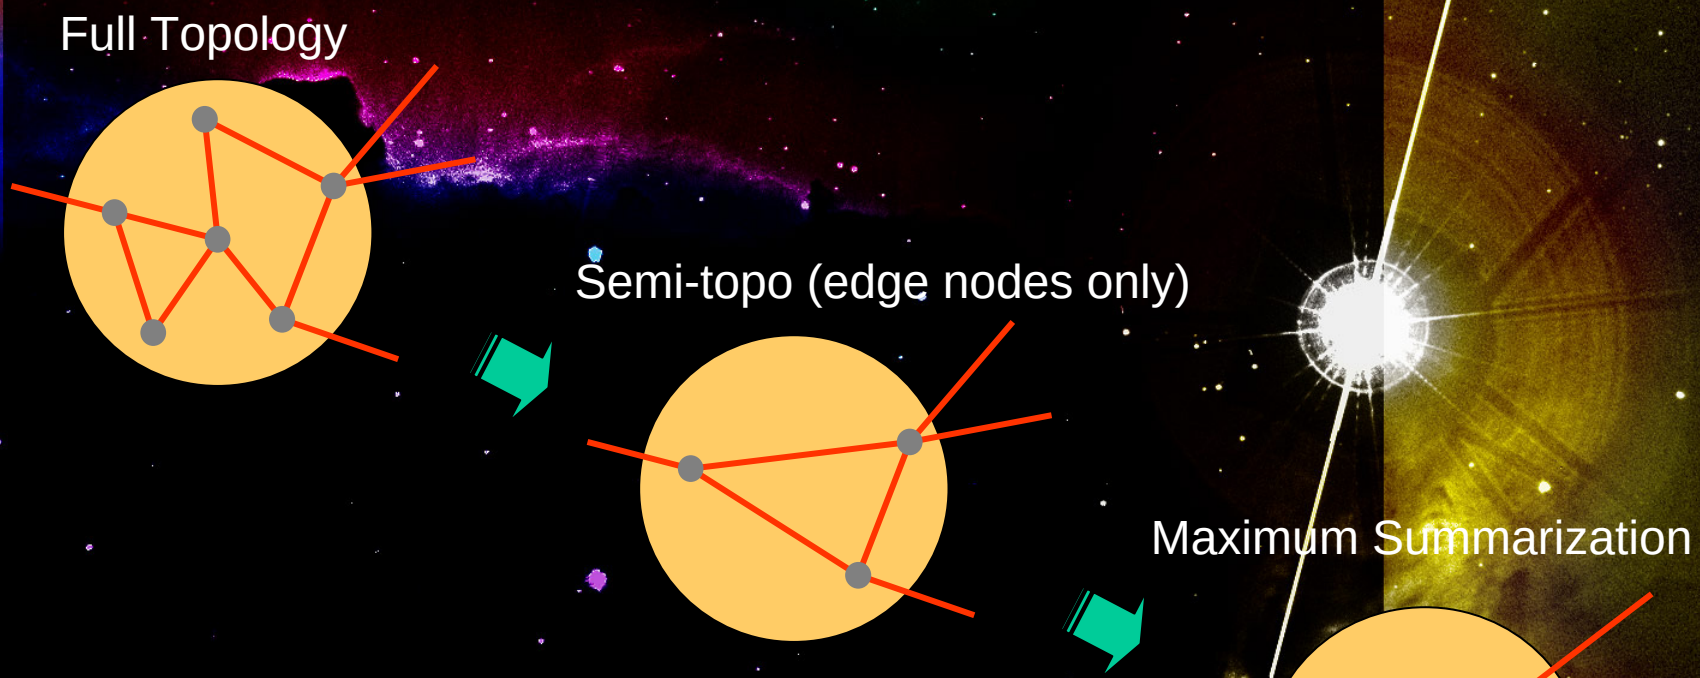

- User defined summarization level maintains privacy
-Summarization impacts optimal PC but allows the domain to choose (and reserve) an internal path

Current DRAGON deployment status:

- DRAGON (in Washington metro area is fully operational
  - L2SC VLSR + interdomain NARB operational
  - ROADMs deployed – wave layer constantly growing and in flux (LSC working, interoperability testing in progress)
  - CSFPF PCE continually being refined
    - Being extended to include Sonet VCAT/LCAS/GFP capabilities
    - Plan to support Ciena CD, Nortel HDX, and Alcatel
    - Initial multi-layer TDM+L2SC this summer (before ONT3/GLIF)
- HOPI has deployed VLSR + NARB
  - Operational since fall 05
- NetherLight is hosting a VLSR + NARB in Amsterdam
  - Operational as of April 2006
  - Peers with HOPI in Chicago via transAtlantic 10G link
- NorthernLight VLSR is in place at KTH – integration with will happen over this summer
- Univ of Manchester has VLSR
- [Hope & expect] Tokyo to have VLSR active by ONT3

DRAGON deployments:
Operational contiguous GMPLS L2SC dynamic reach:

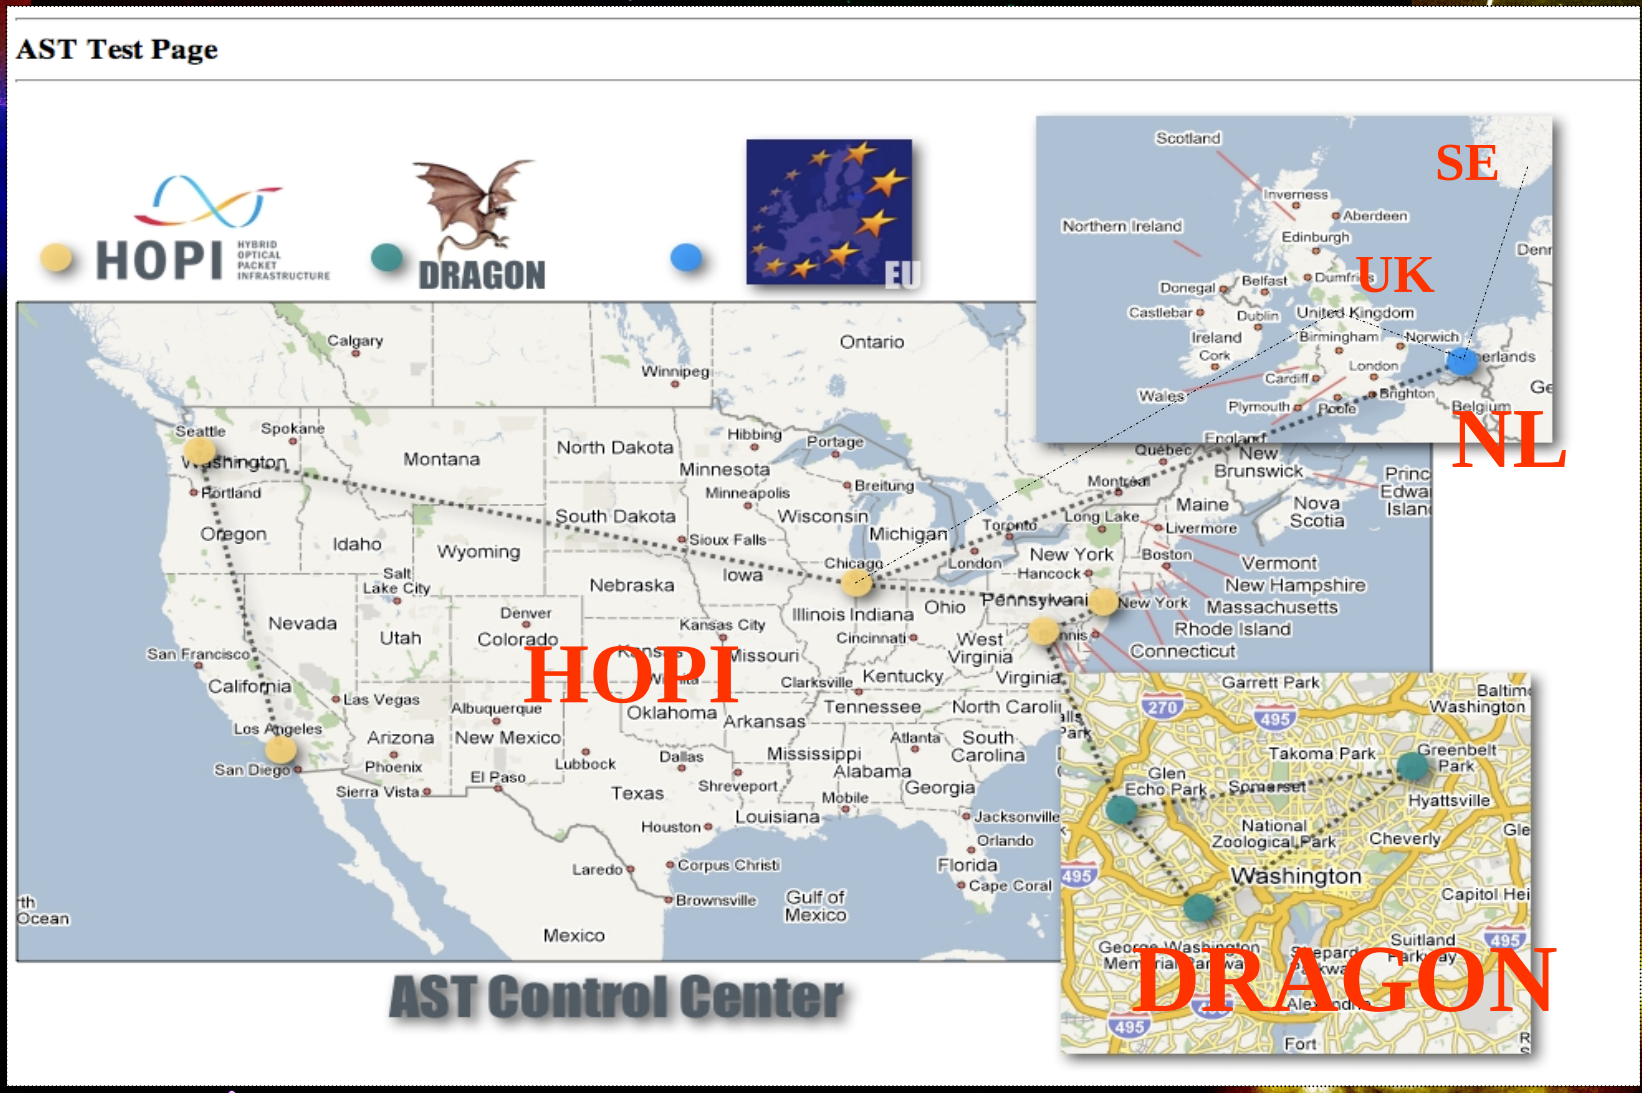

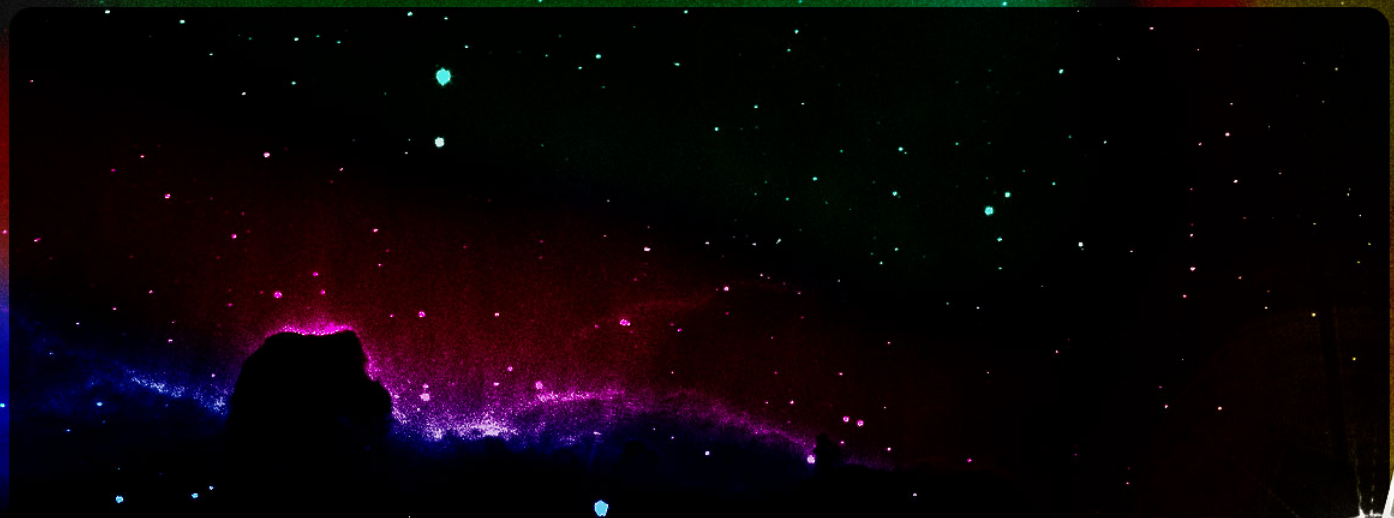

So how is this relevant to Application Controlled Networks?

What are "Application Controlled Networks" ?

- Emerging high performance applications and global collaboration now require predictable and reliable network performance
- Network service providers (both commercial and R&E) are now providing dedicated network resources in the form of waves/circuits/VPNs as part of their service portfolio

- Experimental dynamic networks are now providing applications and end users with the ability to request dedicated network resources that create a "custom" network environment:
  - May be specific capacity or performance between locations…
  - May be particular time that the resource is required, and may be transient
  - May be constrained as to whom is allowed to use it..

- *The user/application essentially gets their own service specific network…*
  - Typically IP is used within the ASN
  - Typically the nodal components of the ASN have separate access to the Internet, the ports/interfaces used here are complementary

An Example:

- Electronic Very Long Baseline Interferometry (E-VLBI)
  - Radio astronomy community shares their resources…An international collaborative effort:
  - 25 to 30 antennae
  - 10 to 15 correlator sites (special DSP hardware)
    - Distributed software correlators are being developed (work being done at JIVE and UvA)
  - Coordination of telescopes to record rf noise from specific celestial targets simultaneously
    - Record 20 minutes to several hours of data at a time
    - Raw data rate of ~512 Mbps per telescope (2006) to 4Gbs/telescope (2010 timeframe) (see the E-MERLIN and JIVE)
  - Real-time or nearRT processing desired for portions of the data, with post processing for most data

E-science Application: Very Long Baseline Interferometry "E-VLBI"

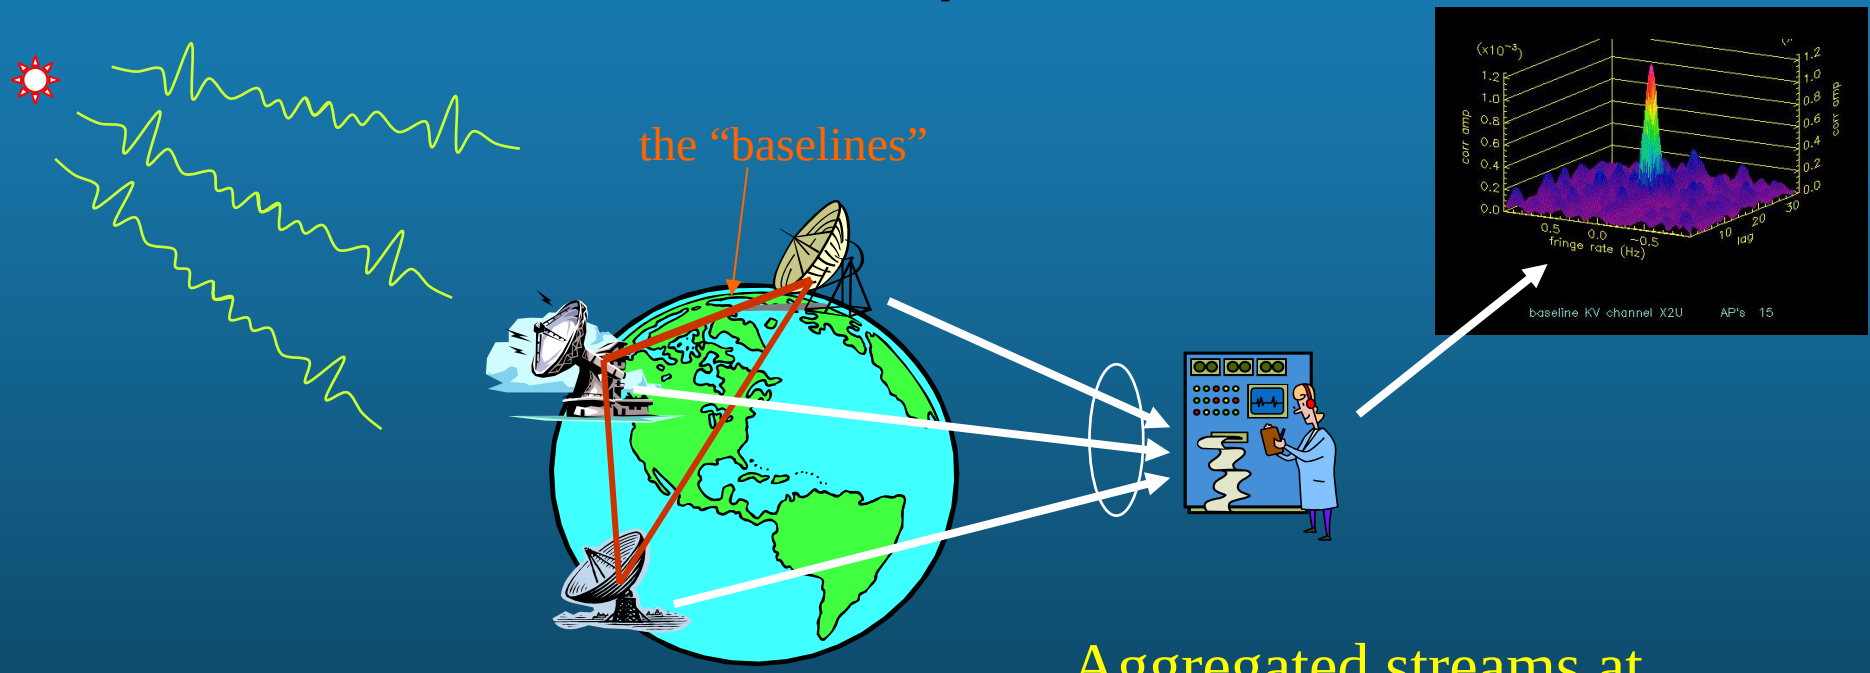

Radio Telescopes
2005 = 512 Mbs
2007 = 2 Gbs
2009 > 4+ Gbs

Aggregated streams at correlator:
2005 > 2 Gbs
2007 ~ 10 Gbs to 20+ Gbs
2009 > 20 Gbs to 40+ Gbs

The "Application Specific Topology"

Logical e-VLBI Topology:

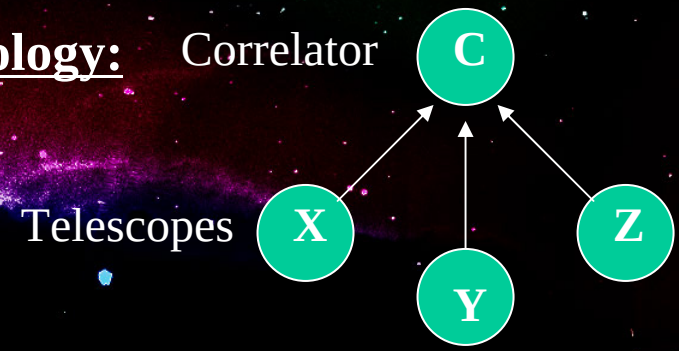

Physical Instantiations of the Application Specific Topology

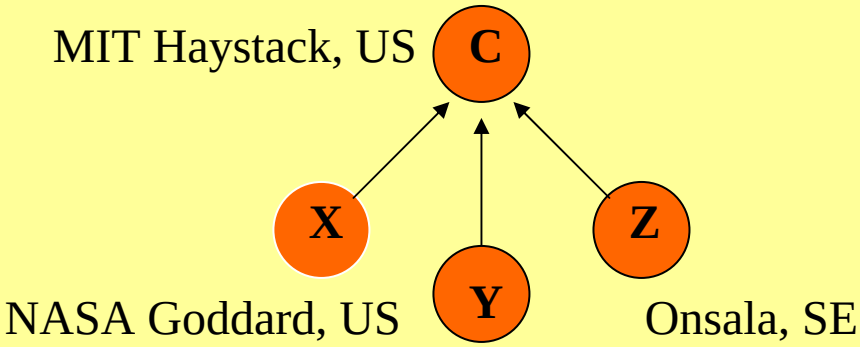

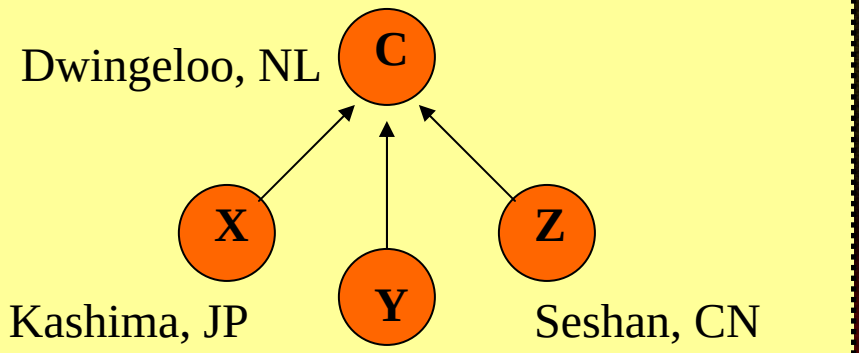

A "VLBI" Application Specific Network

- Telescopes connect to intermediate realtime storage/spooling facilities
  - These storage facilities may be a) at the telescope, b) at the correlator, or c) somewhere else logistically useful.

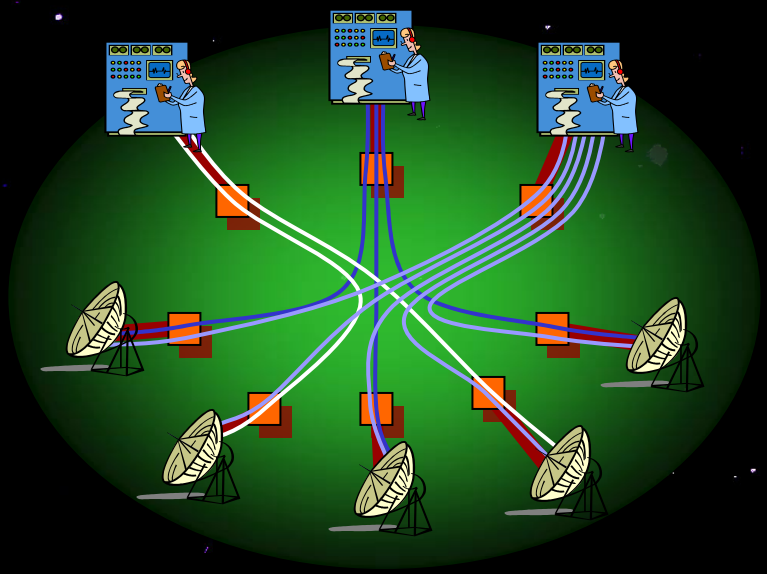

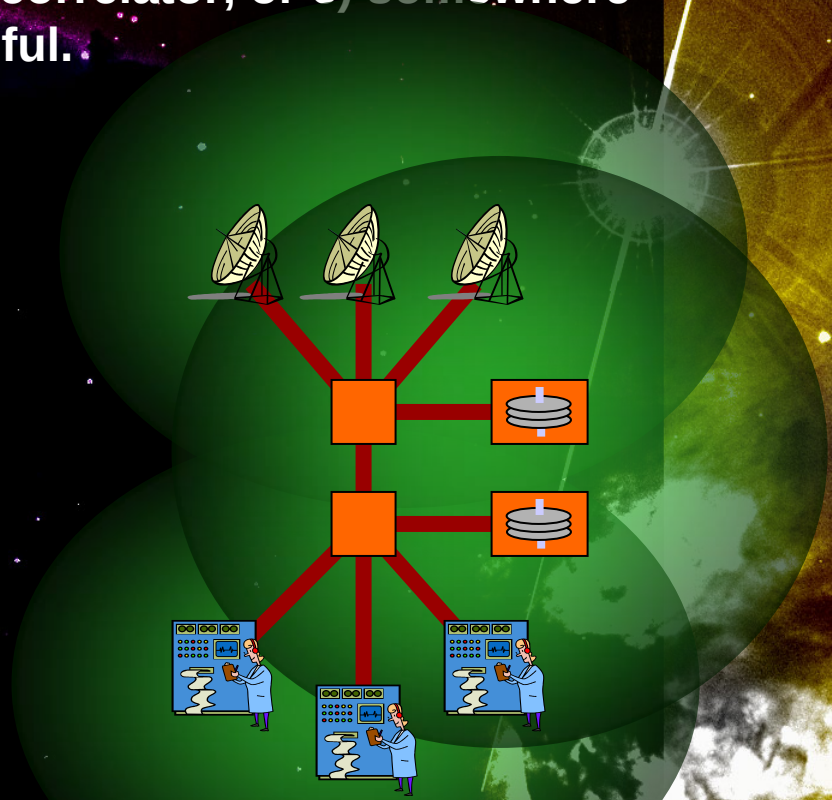

Other Application Controlled Networks

- Bulk Data Transport Services
  - Designed to make file transfer work well even when the end systems are not tuned for TCP over long fat pipes
    - TCP sessions can be intercepted (upon user's request) by Generic Session Layer gateway
    - High performance well engineered links, tuned TCP stacks, and TCP proxy processing exist in the GSL gateways distributed around the world
  - GSL gateways know of each other and construct an internal mesh of high speed transport links
    - End systems hosts talk to local gateways and vice versa

Tha “Black Cloud” project:

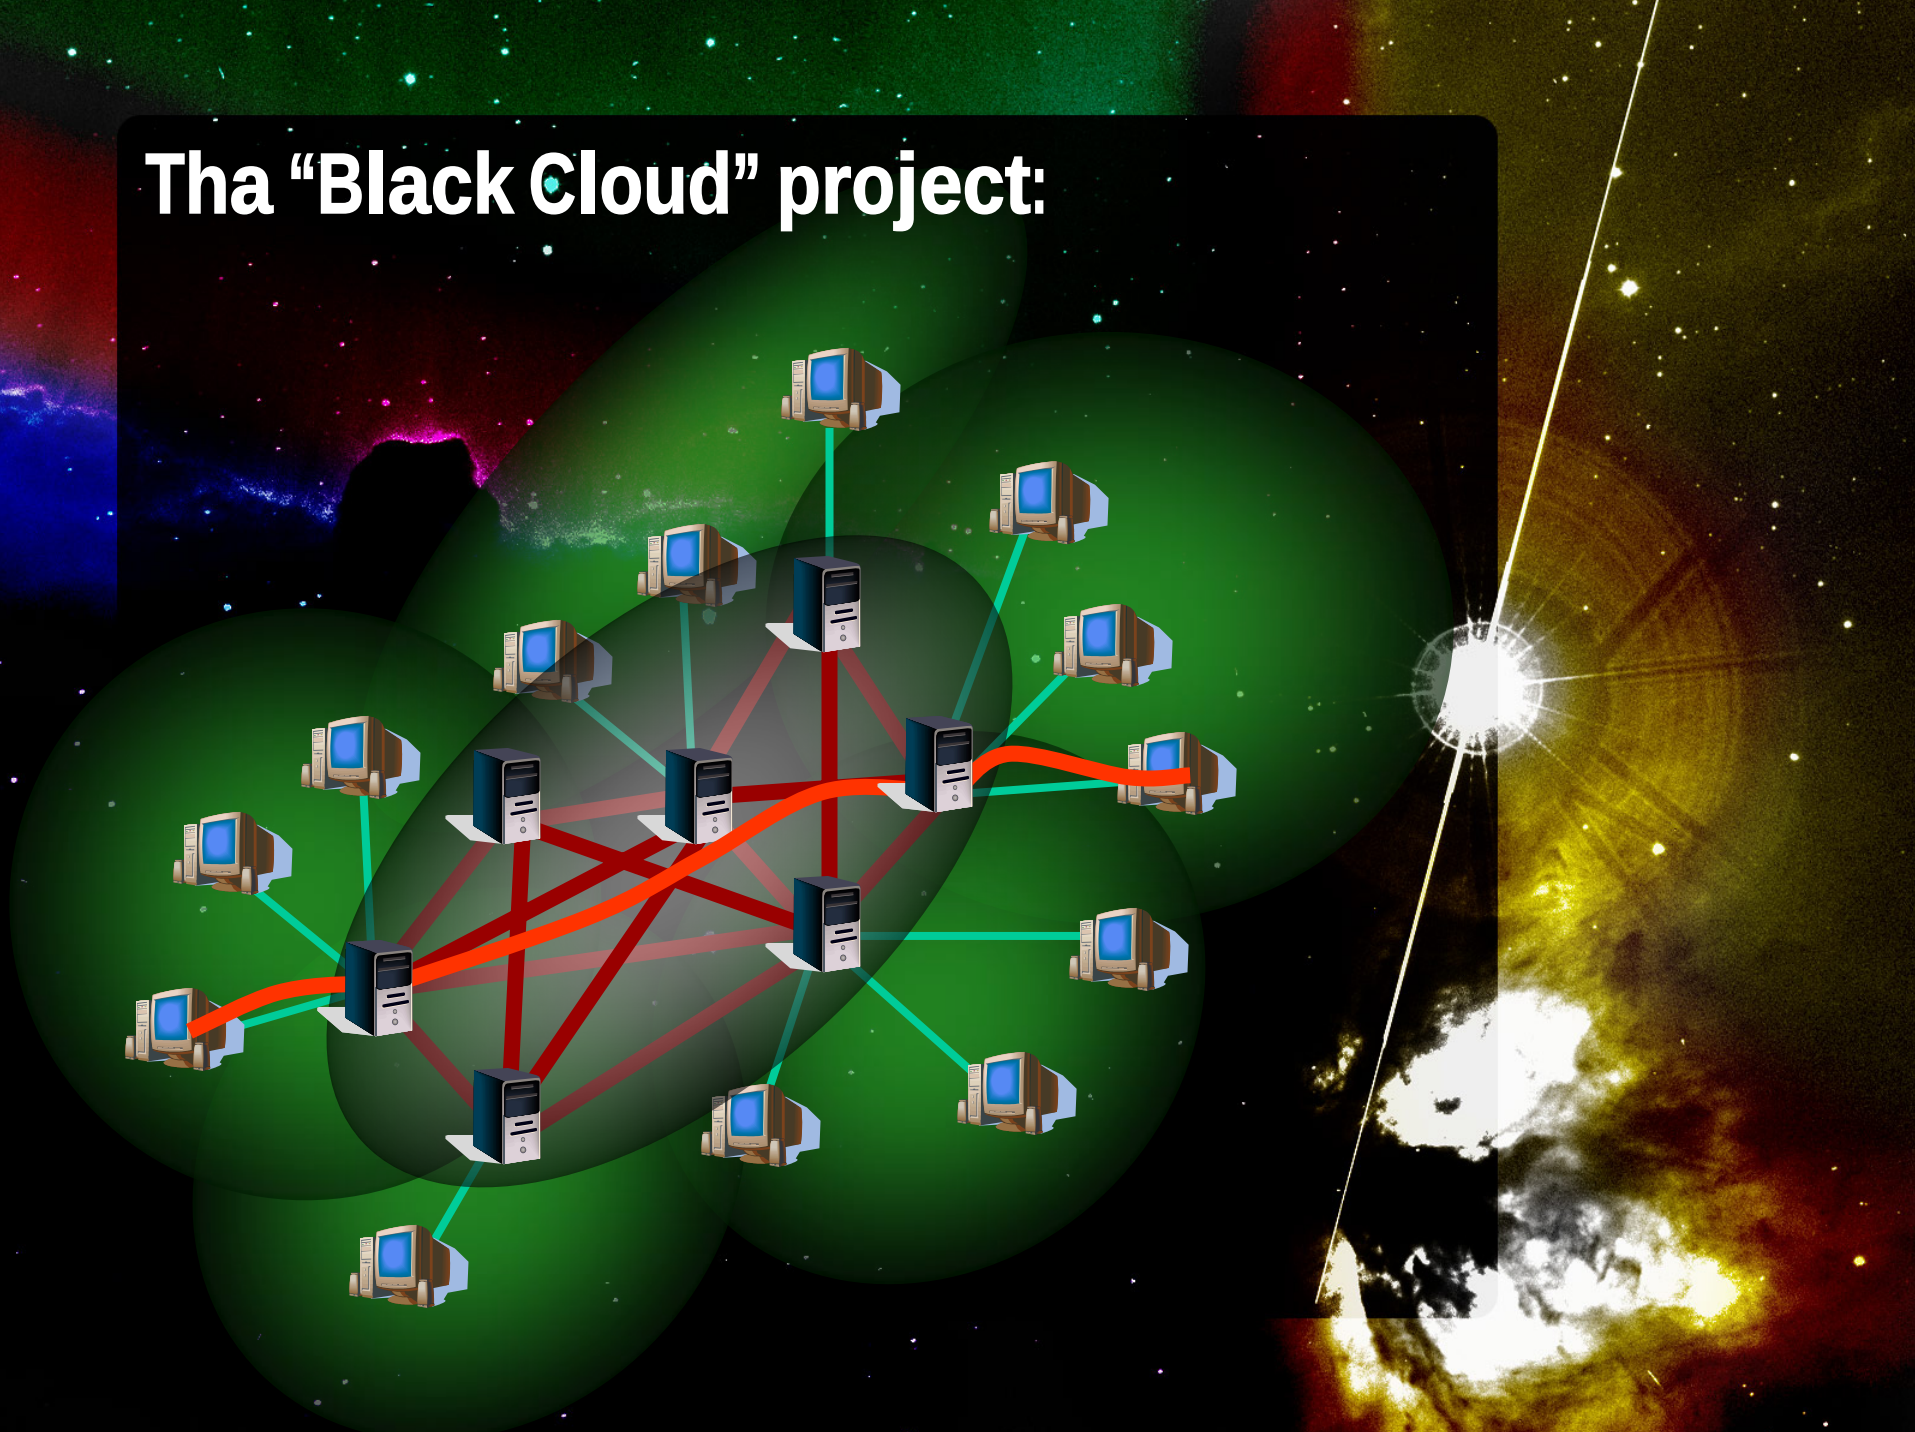

Global e-VLBI

iGrid 2005

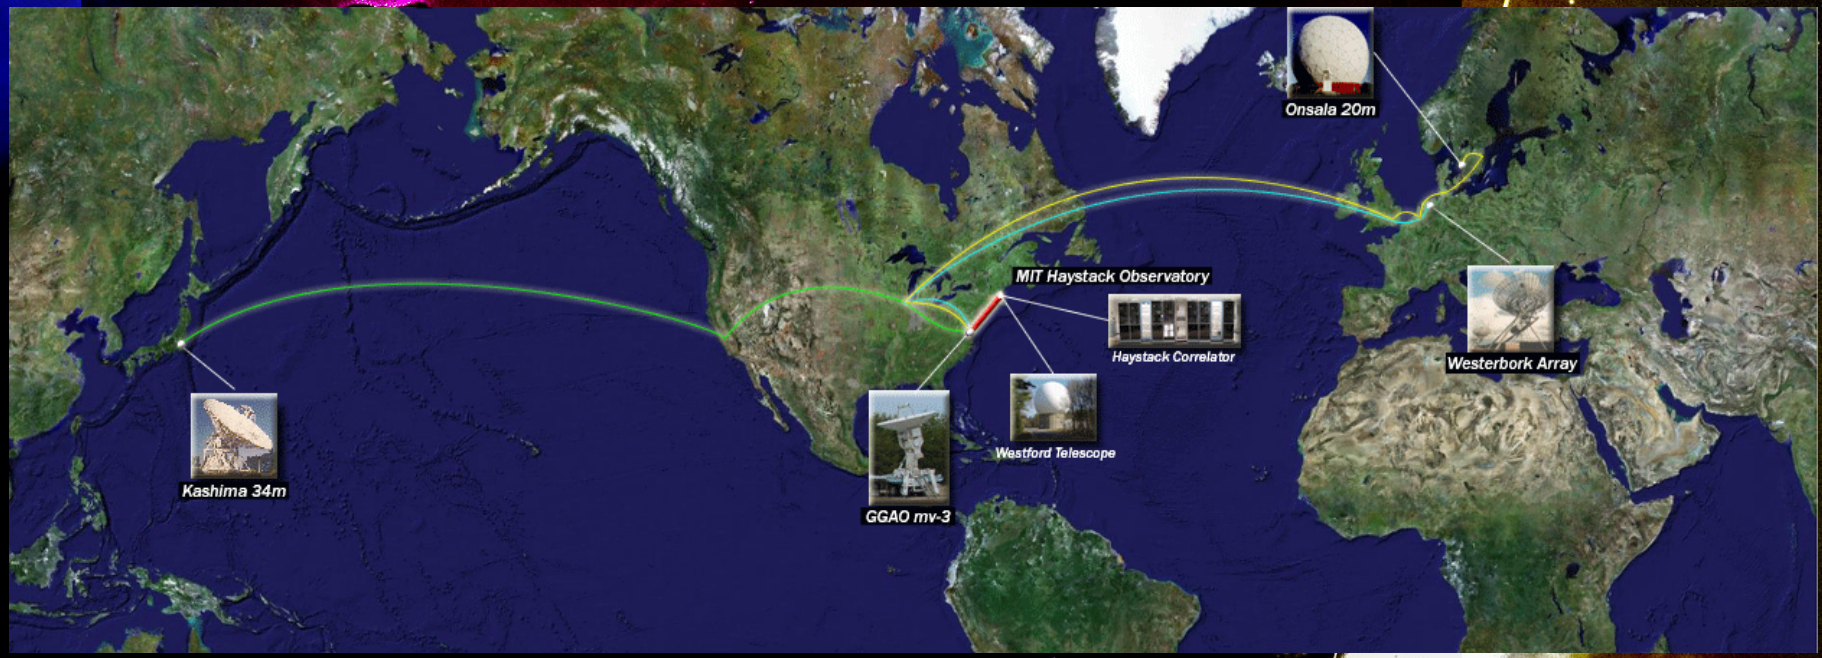

The 30,000 foot view:
HOPI SC05 Demo:

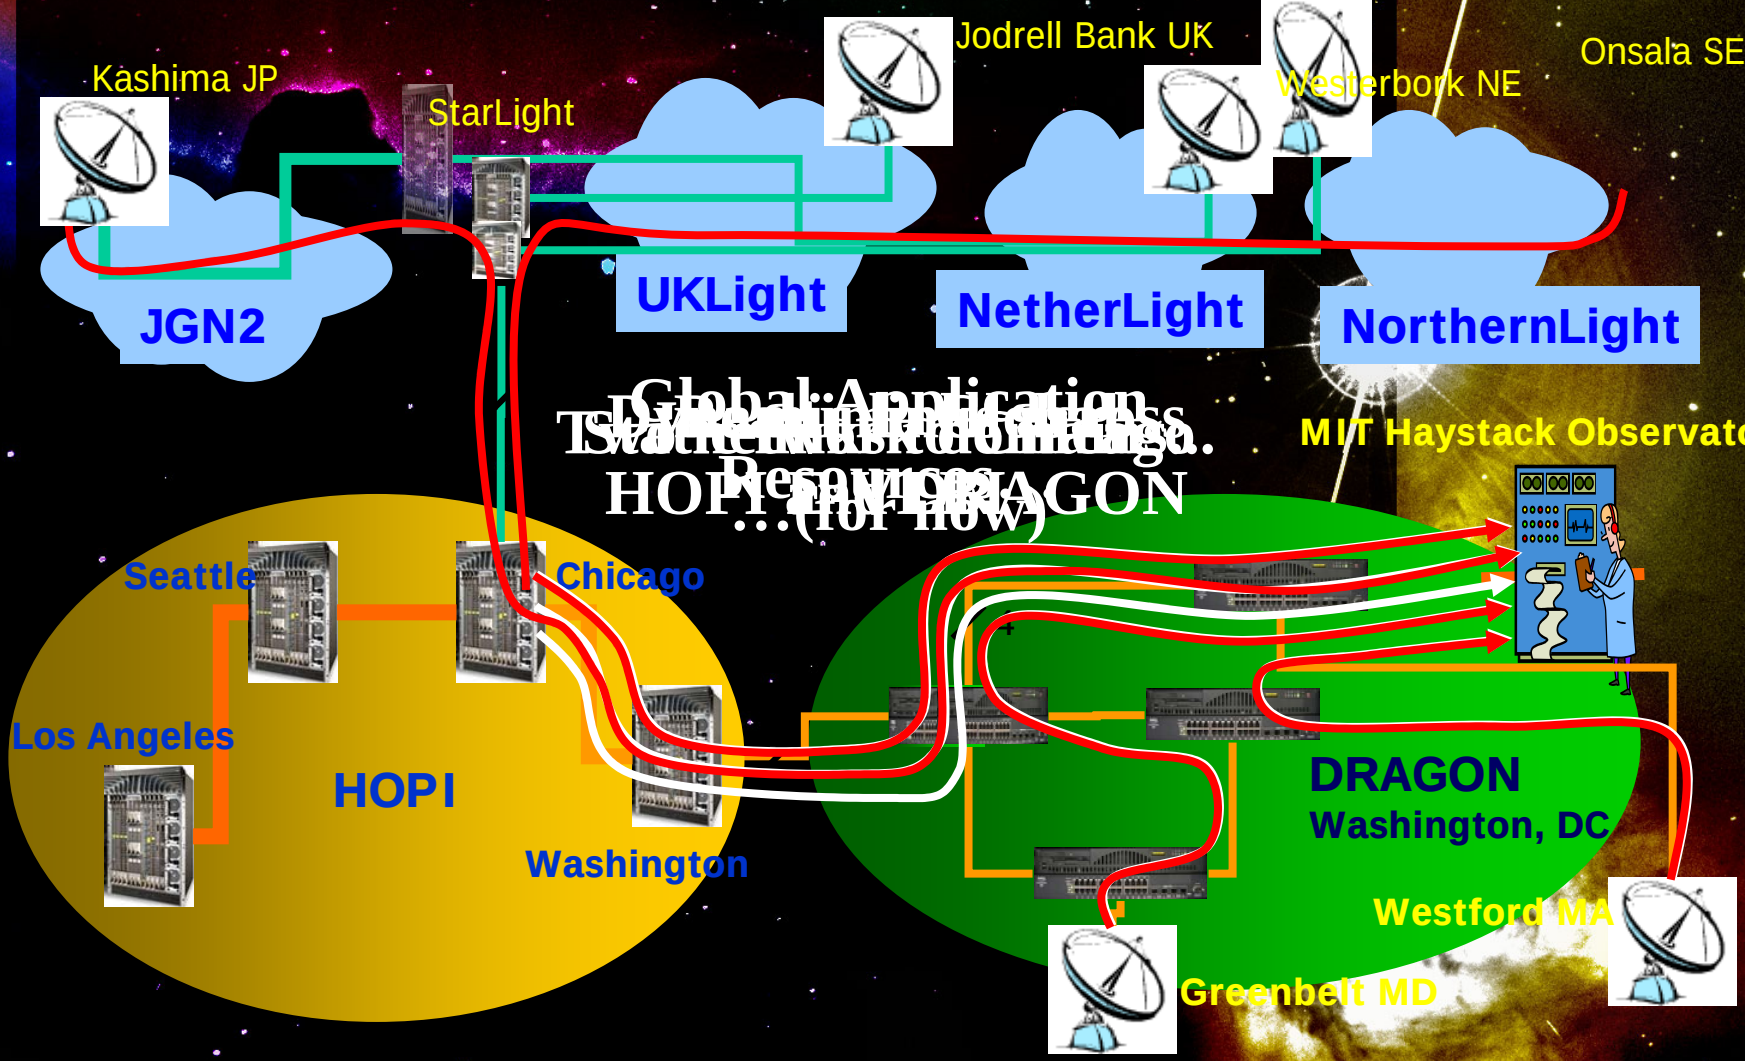

Other Application Controlled Networks

- Video Services
  - Digital video, HD video,
  - Video requires very stringent performance requirements –
    - Compression schemes such as MPEG are extremely sensitive to loss in the network, so engineering long distance video links to eliminate jitter and buffering can reduce loss
    - Compression adds latency, so uncompressed streaming video can significantly improve human factors, but uncompressed requires signifiantly higher bandwdith and performance
  - Solution: develop video gateways/servers and protocols that know of each other around the world. These servers request specific performance requirements of the network anwith ability to request dedicated

Low Latency High Definition Collaboration

LDK-6000

SMPTE-292

e/o

d/a

proxy

NARB

DRAGON Network

proxy

o/e

d/a

e/o

NARB

SMPTE-292

- End-to-end native SMPTE 292M transport
- Media devices are directly integrated into the DRAGON environment via proxy hosts
  - Register the media device (camera, display, …)
  - Sink and source signaling protocols
  - Provide Authentication, authorization and accounting.

HD Collaborative “Video Area Network”

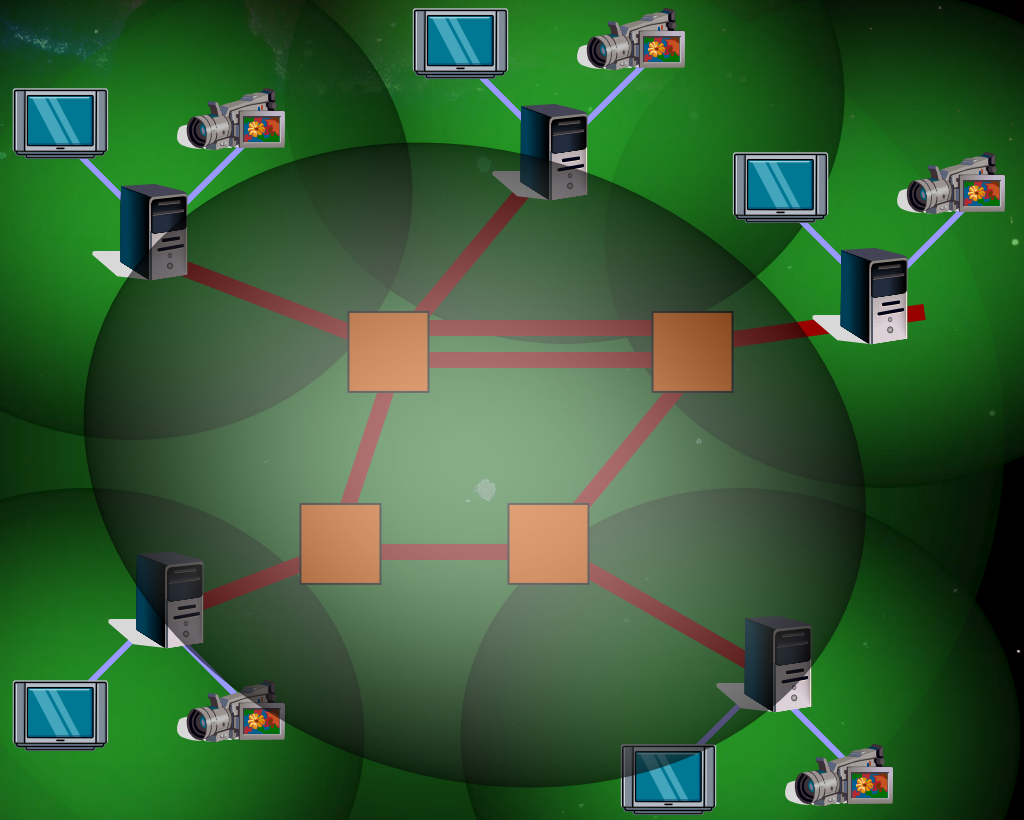

Application Specific Topologies "AST"

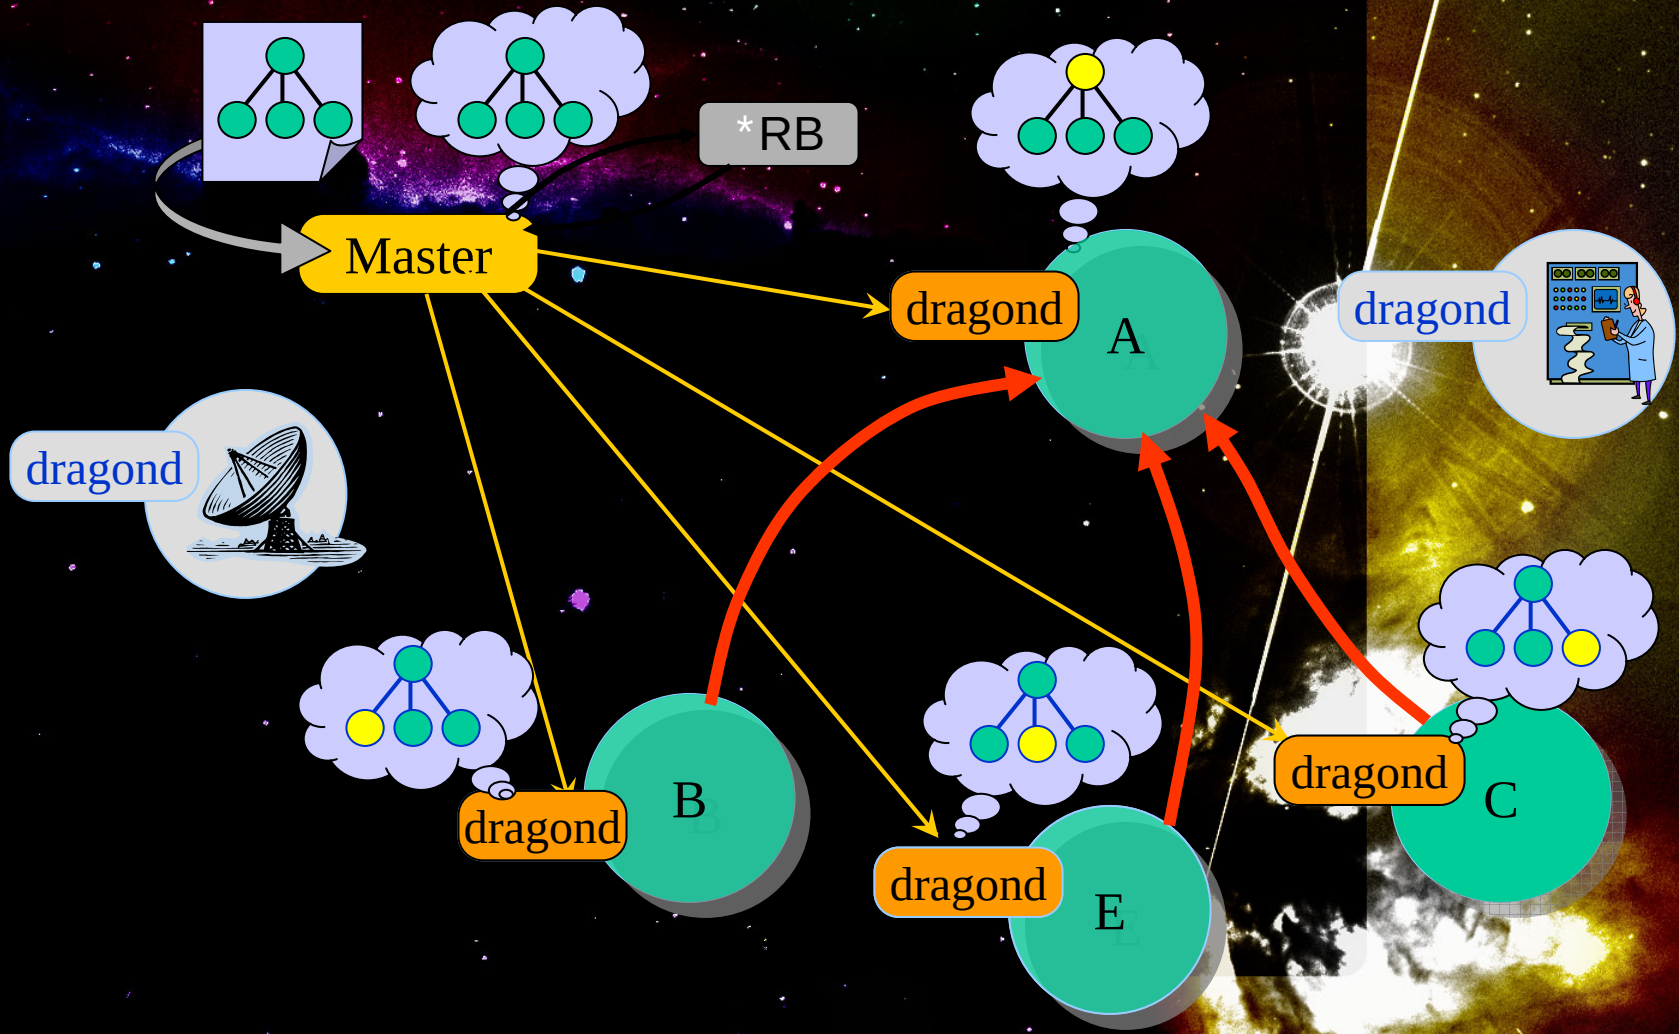

Application Specific Topologies using XML

```
<topology>
    <resource>
        <resource_type> eVLBI.Mark5a          </resource_type>
        <name>     Haystack.muk1              </name>
        <ip_addr>  muk1.haystack.mit.edu      </ip_addr>
        <te_addr>  muk1-ge0.haystack.mit.edu  </te_addr>
        <appl>     /usr/local/evlbi_script    </appl>
    </resource>
    <resource>
        <resource_type> eVLBI.Mark5a          </resource_type>
        <name>     Westford1                  </name>
        <ip_addr>  wstf.haystack.mit.edu      </ip_addr>
        <te_addr>  wstf-ge0.haystack.mit.edu  </te_addr>
        <appl>     /usr/local/evlbi_script    </appl>
    </resource>
    <resource>
        <resource_type> EtherPipeBasic        </resource_type>
        <src>           Haystack.muk1         </src>
        <dest>          Westford.muk1         </dest>
        <datarate>      1 Gbs                 </datarate>
        </resource>
</topology>
```

A

B

C

The AST Process

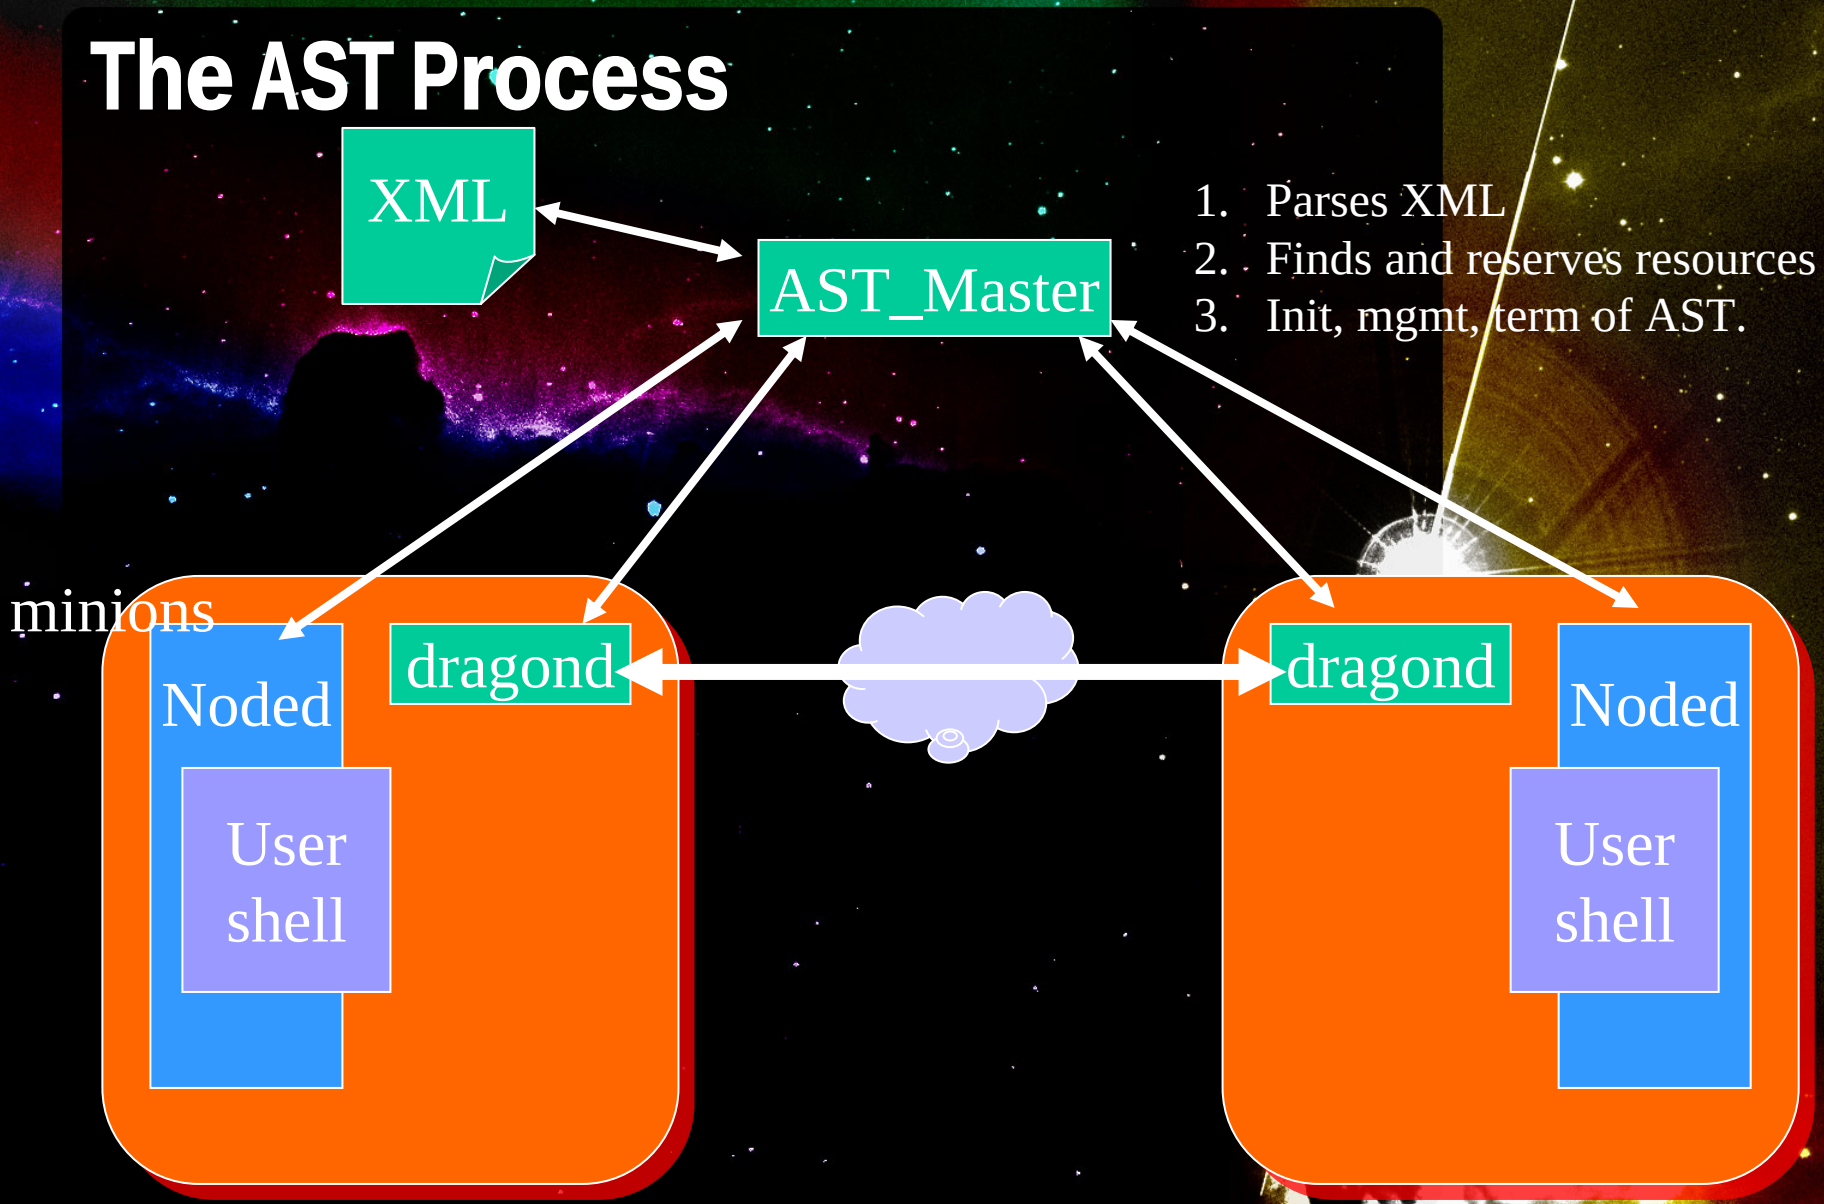

The AST Protocol

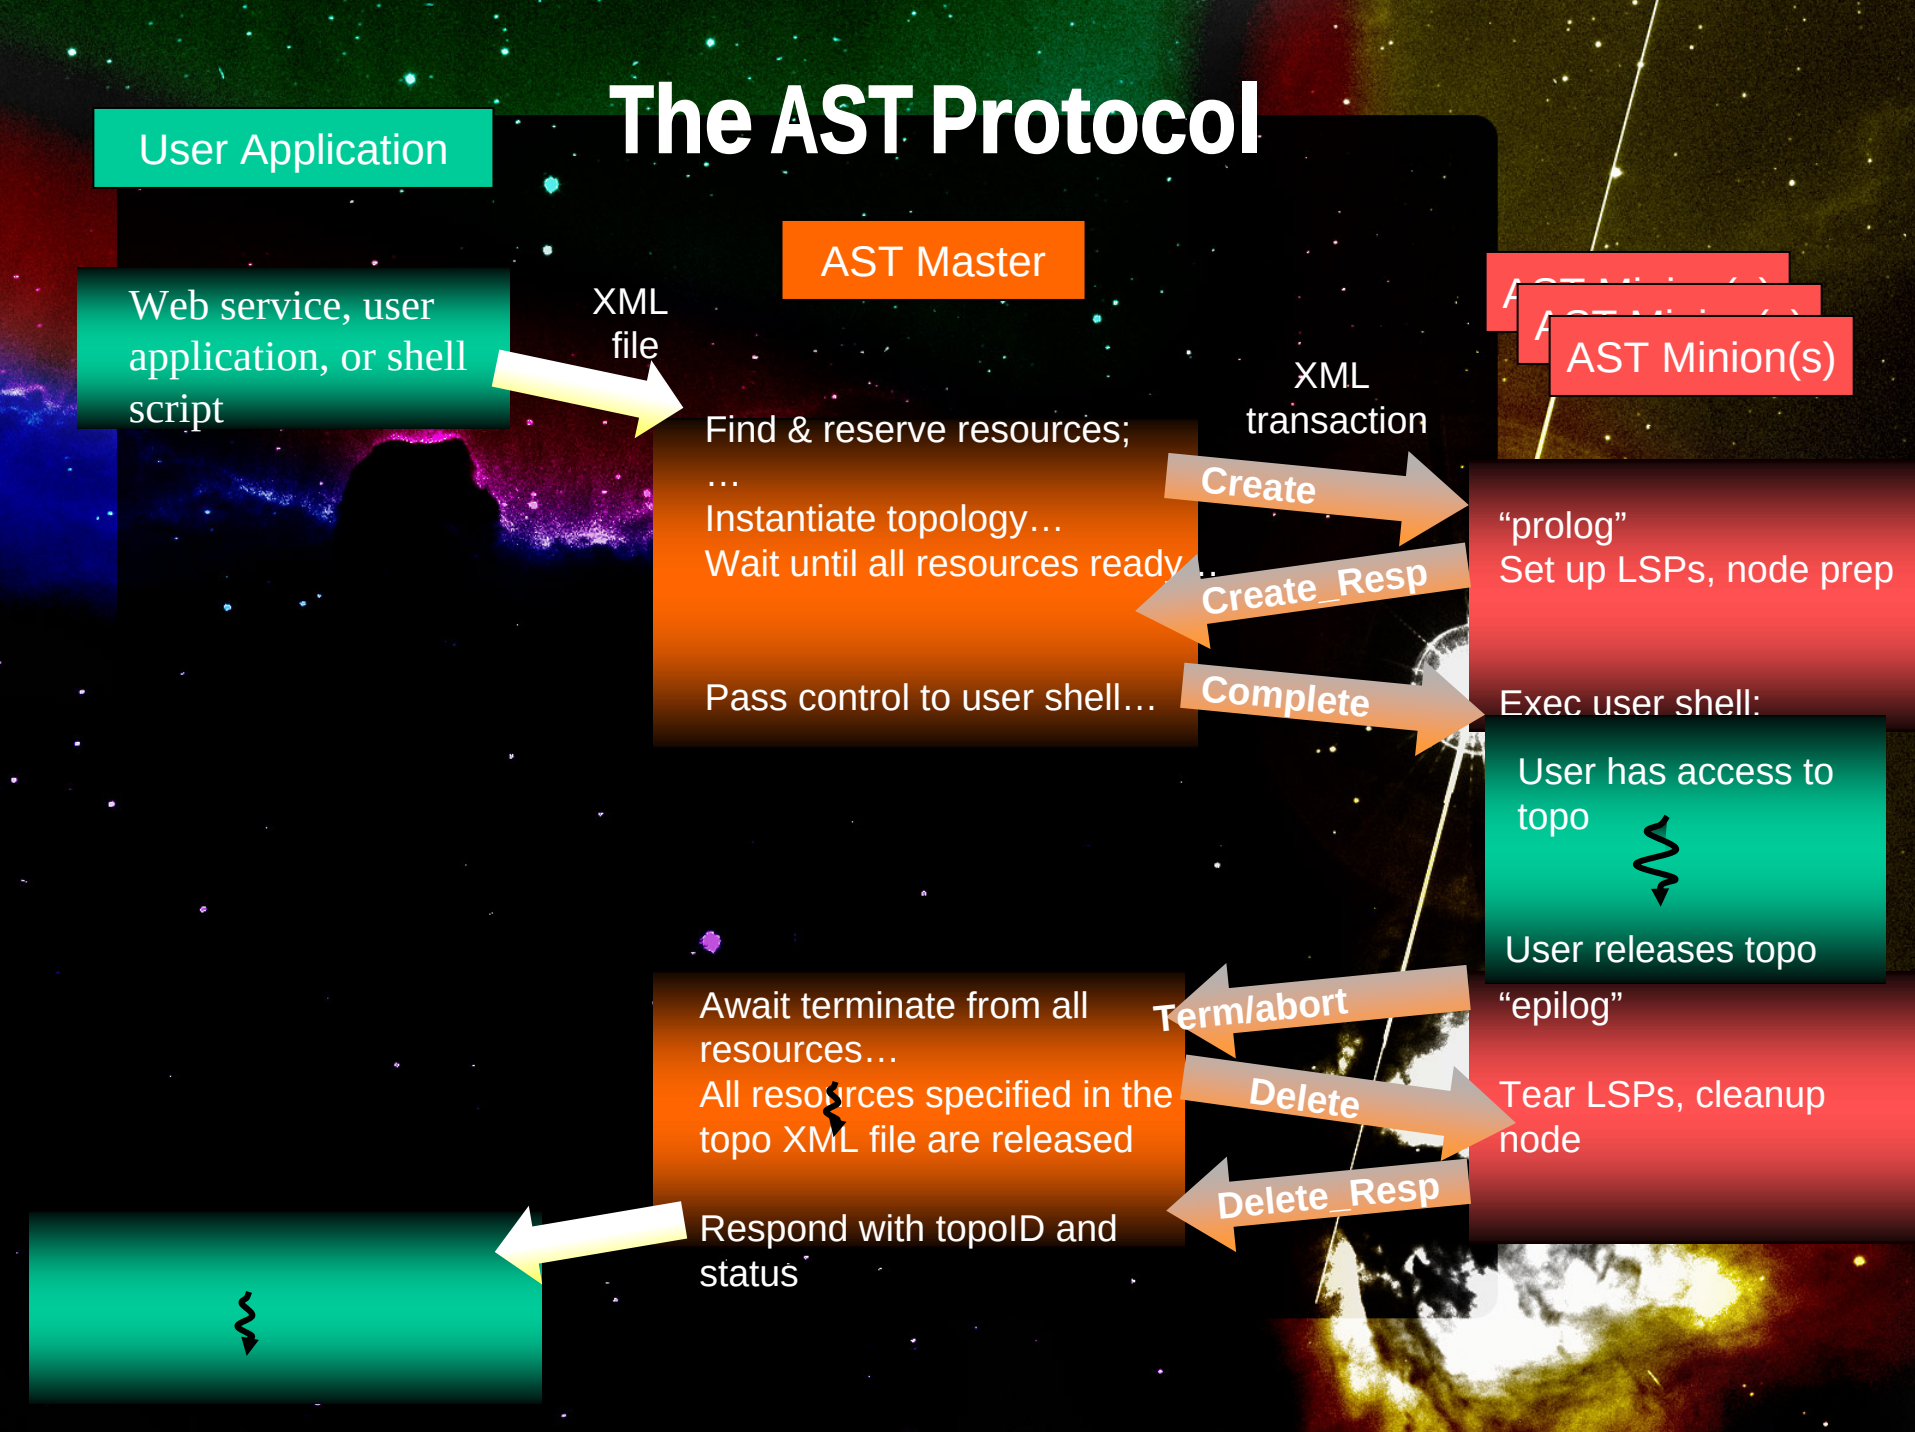

Applications Specific Topologies

- Live demonstration at Internet2 Spring Member Meeting (April 2006, Washington DC)
  - See www.internet2.edu for webcast of "HOPI update" presentation. (See me after this talk for personal live demo (☺))
- Set up global multi-link topologies
  - Less than 30 seconds (!)

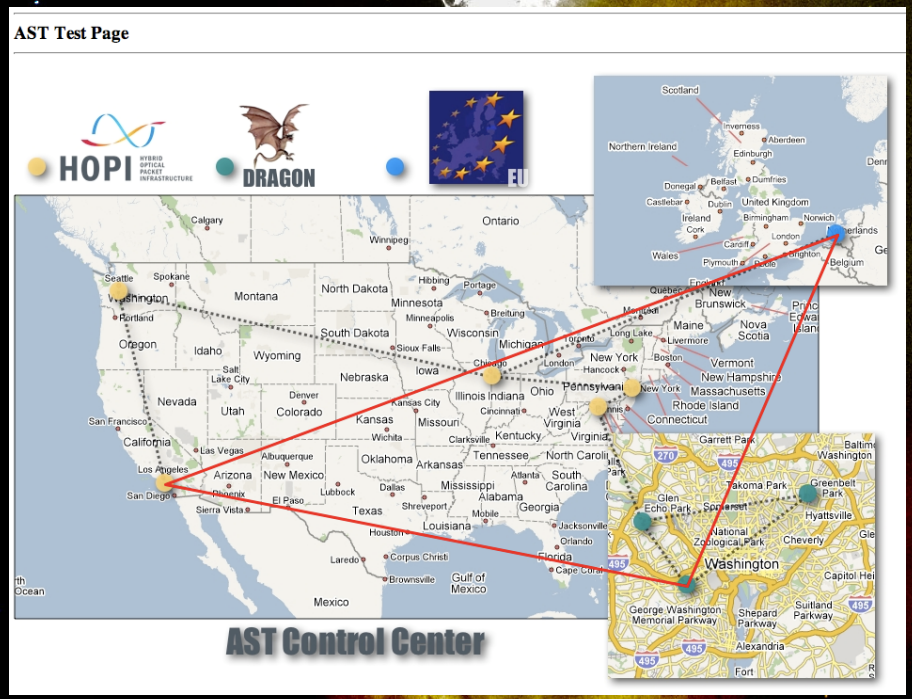

Summary

- Application Specific Topologies (ASTs) (Service Specific Networks, Service Oriented Networks, etc)
  - Network topologies established to provide or optimize specific services, or to address the needs of a group of collaborating people or organizations.
  - User defined topologies set up en-masse
  - These will be a substantial new capability in next gen networks that the user community will expect.

- Areas for further research and development:
  - Interoperability with a) multi-layer functions, and b) vendor stacks
  - Development and integration of Network Resource Brokers into the global grid environment
  - Service definitions to define service capabilities of emerging networks – Ethernet? Sonet/SDH? Packet? Heterogeneous?
  - Advanced scheduling and AAA – and integration with 2-stage reservation and provisioning architecture
  - AST protocol development – and reconciliation with related architectures such as WS, ACNs, etc.
  - What will we do to carry these ideas into the Tbs regimen?

- **Thank You!**
  - Jerry Sobieski jerrys (at) maxgigapop.net
  - **> dragon.maxgigapop.net**
    - **Free software**
    - **Free support**
    - **Free global dynamic lightpaths (for now☺)**