and

IEEE 802.1q VLAN-based R&E Network Exchange Services

For the JET Roadmap Workshop
Wednesday, 14 April 2004

Bill Nickless

nickless@mcs.anl.gov

(This week!)

WKN 20040414

Chicago Situation in mid-2001

- Ameritech AADS NAP (STARTAP service provider) limits:
  - No services beyond OC-12 ATM
  - Hard to interface with Ameritech (Bell Operating Company)
  - No co-location space available
- I-WIRE fiber network funded
  - Customer owned (capitalize bandwidth)
- OMNInet research project coming online
  - Multiple 10GigE channels at several Chicago sites
- StarLight Facility:
  - Carrier-neutral co-location focused on R&E network activities
  - I-WIRE, OMNInet, SURFNet OC-192 from Europe, etc.
  - Opportunity to implement circuit, packet, and/or frame exchange

Chicago's international R&E network exchange point must move to the next level

STARLIGHT℠
The Optical STAR TAP™

Beyond the AADS NAP:
(What They Got Right)

- "Bring your ATM Cells to Chicago": meet many interesting and valuable peer networks (at DS-3, OC-3, or OC-12 speeds)
- TCP-friendly ATM switches (deep queues)
- Point-to-point connectivity between participants (ATM PVC mesh)
  - Customer controls to whom you send and receive
  - No Acceptable Use Policy (AUP)
  - No Technology Policy (IPv4/6, DECNET, IPX, AAL3 Video, etc. all supported transparently)

IEEE 802.1q VLANs at StarLight:
Update the AADS NAP strategy

- "Bring your Ethernet Frames to Chicago" to meet interesting and valuable peer networks (at 10GigE, 1GigE, or Fast Ethernet speeds)
- Point-to-point-ish connectivity for participants
  - No Acceptable Use Policy (AUP)
  - No Technology Policy (IPv4, IPv6, DECNET, IPX?)
  - Avoids ugly IPv4 inter-domain multicast issues
  - Protect participants from each other
- TCP Friendly
  - "Speed bumps": factors of 10 rather than factors of 4

Bilateral IEEE 802.1q VLANs

- MTU is negotiated between individual customer peering sessions
  - Up to 9216 bytes per frame at StarLight
- Limited IEEE 802 broadcast domain
  - IP Multicast MSDP/M-BGP/PIM-SM ugliness
  - Constrain flooding (accidents or DOS attacks)
- StarLight does not manage Layer 3 addresses
  - No common IPv4 peering subnet
  - StarLight avoids critical path for IPv6 (or whatever)

STAR TAP℠ and STARLIGHT℠
The Optical STAR TAP™

Services Description (Mid 2002)

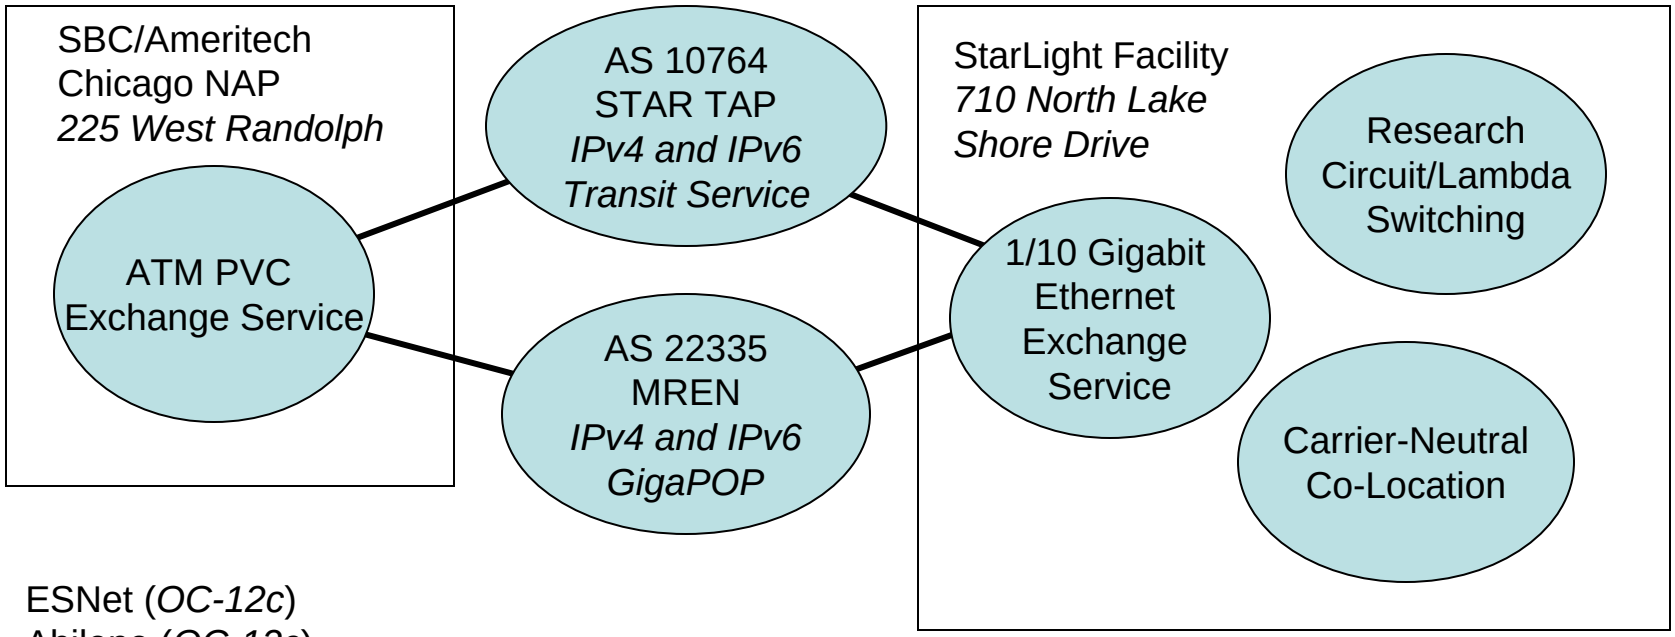

ESNet (*OC-12c*)
Abilene (*OC-12c*)
vBNS+ (*OC-12c*)
Fermilab (OC-3c)
Argonne (*OC-12c*)
CA*Net3 (*OC-12c for 1 more day*)
Illinois Century Network (*OC-12c*)

Abilene (*2xGigE, 10GigE soon*)
TeraGrid (*10GigE soon*)
SURFNet (*2xOC-12c, OC-48c*)
Argonne (*2xGigE*)
Northwestern University (*GigE*)
AMPATH (*OC-3c*)
NaukaNet (OC-3c; *was FASNet*)
ESNet (*GigE*) NREN (*GigE*)

I-WIRE (dark fiber/vFMMR)
TRANSPAC (OC-12c)
NorduNet (OC-3c)
vBNS+ (*GigE soon*)
CA*Net4 Winnipeg (GigE)
CA*Net4 Toronto (GigE)
NCSA (2x*GigE very soon*)
CERN (*OC-12c, OC-48c soon*)

Configuration Example

```
interface GigabitEthernet2/9
   description NISN/NASA
   mtu 9216
   no ip address
   speed nonegotiate
   switchport
   switchport trunk encapsulation dot1q
   switchport trunk allowed vlan 210-213,217-226,231,232
   switchport mode trunk
   switchport nonegotiate

interface GigabitEthernet2/10
   description GEMnet
   mtu 9216
   no ip address
   speed nonegotiate
   switchport
   switchport trunk encapsulation dot1q
   switchport trunk allowed vlan 167-169,231
   switchport mode trunk
   switchport nonegotiate
```

StarLight 802.1q VLAN Ids

30-Jul-02 11:11

MREN 6509 (AS 22335)																		
101	STAR TAP (AS 10764)																	
100	102	SURFNet (AS 1103)																
107	103	104	Abilene (AS 11537)															
				NWU (AS 103)														
	105				AMPATH (VLAN 105 only)													
128		106				6TAP												
108						128	MREN M5 (AS 22335)											
109								Argonne (AS 683)										
113	110	112	111	117				118	NREN (AS 24)									
119	120		121						114	CERN								
124	125	126							115	122	ESNet (AS 293)							
	129								116			NaukaNet						
123	130	134	135			131			133	136	132		CA*net-Winnipeg (AS 6509)					
147	146	145	144			143			142	141	140		X	CA*net-Toronto (AS 6509)				
127															OMNINet			
													137	148		vBNS (AS 145) *[Unknown]*		
138																	EVL	

Enhancing The Model

- Getting past the 802.1q 4096 VLAN limit: Closer to the “Full PVC Mesh” model
  - Rewrite tags in the switches? (Not likely)
  - Juniper Circuit Cross Connect (CCC)?
  - “Martini Draft” Layer 2 MPLS VPNs
- Distributed Exchange Services
  - Multiple IEEE 802 switches with 802.1q trunks
  - “Martini Draft” Layer 2 MPLS VPN services