NetApp™

Go further, faster

NetApp's Transition and Use of Global MPLS IP VPN

Michael J. Morris

Team Lead and Network Architect

NetApp – The Company

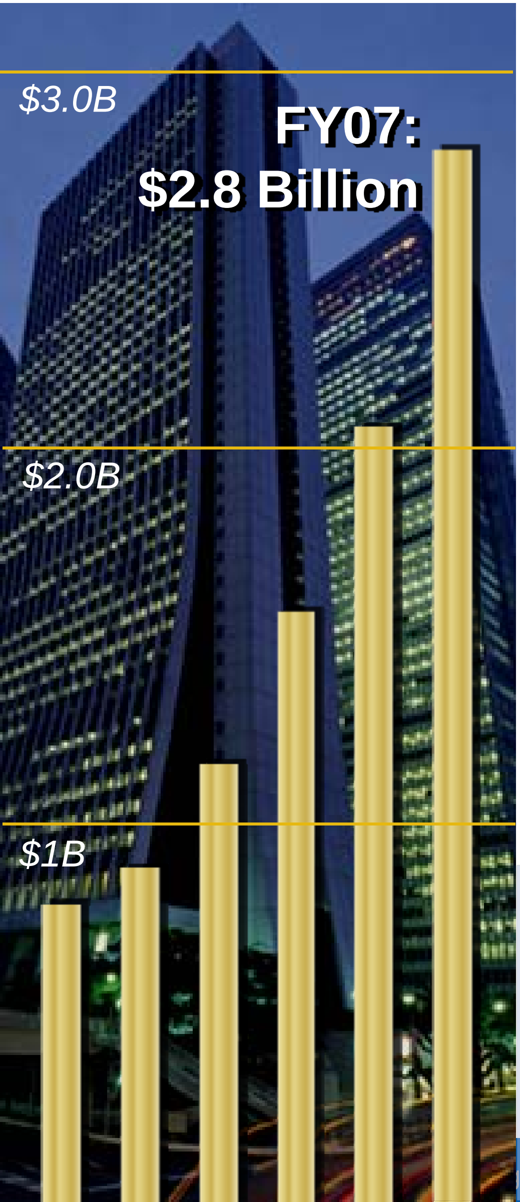

- Worldwide, enterprise customers
- Fastest growing storage company
  - Outpacing the industry by 3x
- Industry-leading partners
- Comprehensive professional services with global support
- Headquartered in Sunnyvale, CA with major offices in RTP, NC, Amsterdam, and Bangalore, India.

- 6500+ Employees
- Distributed in 138+ countries
- 94,000+ installed systems

- Fortune 1000
- S&P 500
- NASDAQ 100

Why We Built a New Network

- Legacy network issues (performance, scalability, flexibility, capacity)
- Applications
- Data center consolidation
- Demanding users
- Internet centralization and security
- Cost reduction and ROI

NetApp™

Go further, faster

Architecture and Templates

A Documented Architecture

- Standards drive consistency, simplify deployment, and ease operational support.
- Nothing is a standard until it is agreed to and written down.
- Use of wiki pages and MS SharePoint
  - Tracks all changes and reverts to previous states if necessary
  - Configurations, guidelines, and general information captured for all to use.
  - Dynamic and easy to use.
- Our “playbook” - single point of reference for all network standards.
- Architecture Review Board

Templates

- Templates (Visio) provide a standard, repeatable design for all sites.
  - 5 field site tiers with corresponding design templates.
  - Other templates for VPN, DMZ, Data Centers, Load Balancing and IP Telephony.
  - If we do a new design, we turn it into a template for future use.
- Template Format
  - Standard icons and diagram template.
  - High-Level, Physical, Logical, Layer-3 Routing.
  - Future additions (Layer-0, Security).
  - Use in DC Networks.
  - Spreading to other IT groups.

Template Example - Physical

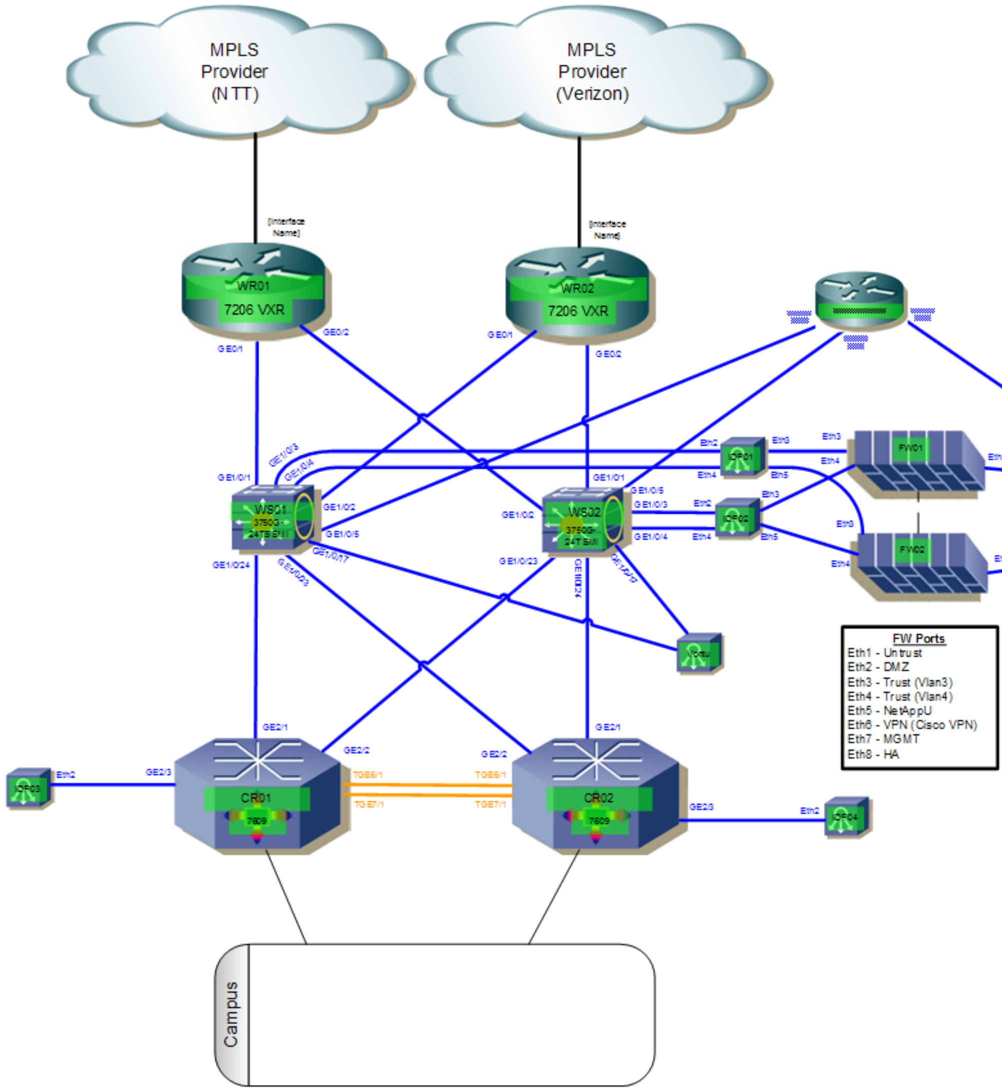

FW Ports

Eth1 - Untrust
Eth2 - DMZ
Eth3 - Trust (Vlan3)
Eth4 - Trust (Vlan4)
Eth5 - NetAppU
Eth6 - VPN (Cisco VPN)
Eth7 - MGMT
Eth8 - HA

Physical Notes

- A single Internet Router (IR) and a single, non-stacked 3750 (IS) can be used in the event only a single Internet circuit is used at the site. There should be two of all other Internet devices.

- WAN Ethernet connections for MPLS should use the NM-2FE2W-V2 card for the WAN connections. All other WAN types should use the appropriate interface card.

IDP Notes

- IDPs are in-line for Internet traffic, not corporate traffic.
- IDPs are in passive mode for internal traffic from the core.
- Separate IDPs are used for Internet and passive mode IDPs. A single IDP can be used for the passive IDPs if the BW from the SPAN off the core routers is not too great. A simple NetFlow analysis can determine if a single IDP will suffice. SecOps will determine if one or two IDPs is needed.
- There is no switch between the IDPs and FWs. This is a direct connection between these devices. All interfaces should be set to auto/auto.
- Tier-1 sites without Internet access will be provided a passive IDP on a case-by-case basis determined by SecOps.
- Vontu is in passive mode for Internet traffic with 'span' from WAN Switches. Traffic to/from the FWs (GE1/0/3-4) are spanned to the Vontu.

NetApp Proprietary Use Pursuant to Company Instructions	Title	Physical Design	Project	Tier 1 Template - WAN	Change Control		Revision	4.0
			Author	Michael J. Morris	Date	1 JUL 07	File Name	Tier1Template-WAN.vsd

Template Example - Logical

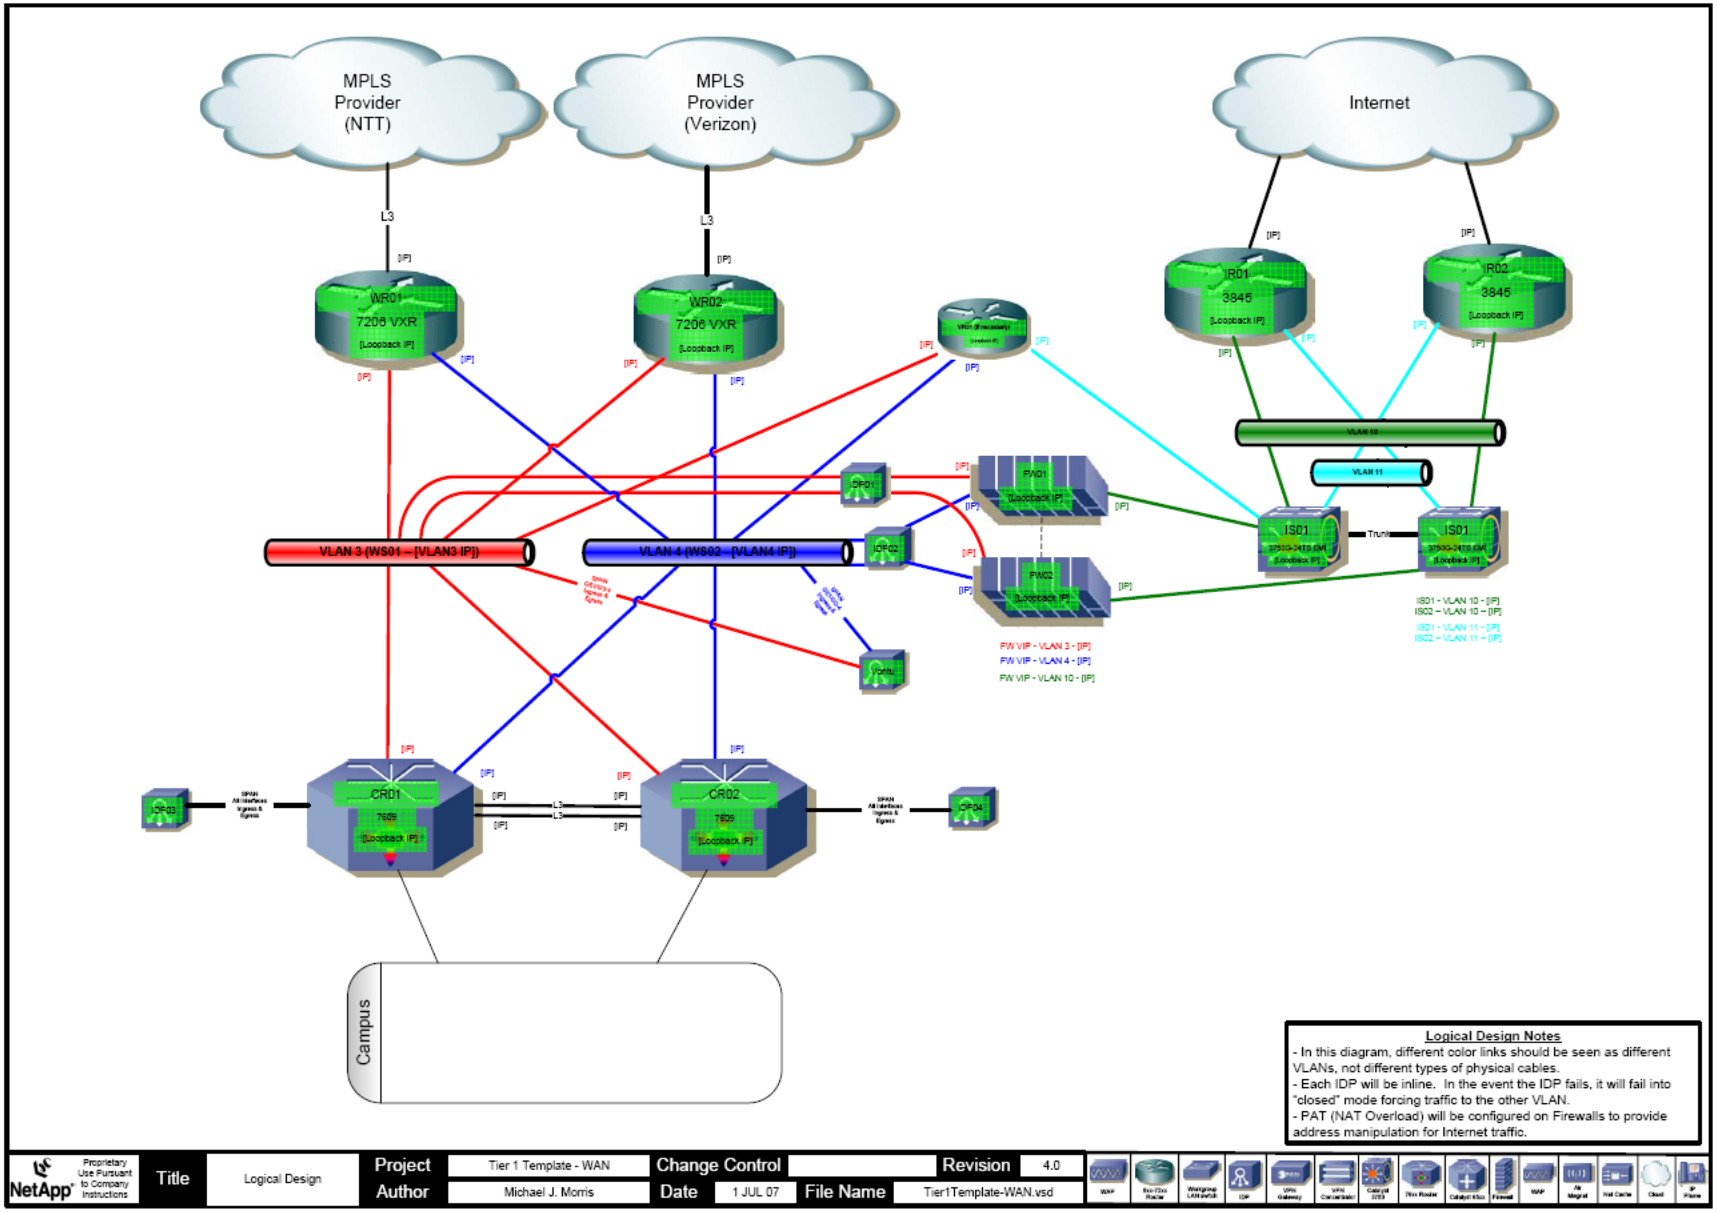

Logical Design Notes

- In this diagram, different color links should be seen as different VLANs, not different types of physical cables.
- Each IDP will be inline. In the event the IDP fails, it will fail into "closed" mode forcing traffic to the other VLAN.
- PAT (NAT Overload) will be configured on Firewalls to provide address manipulation for Internet traffic.

NetApp Proprietary Use Pursuant to Company Instructions	Title	Logical Design	Project	Tier 1 Template - WAN	Change Control		Revision	4.0
			Author	Michael J. Morris	Date	1 JUL 07	File Name	Tier1Template-WAN.vsd

Template Example – Layer-3 Routing

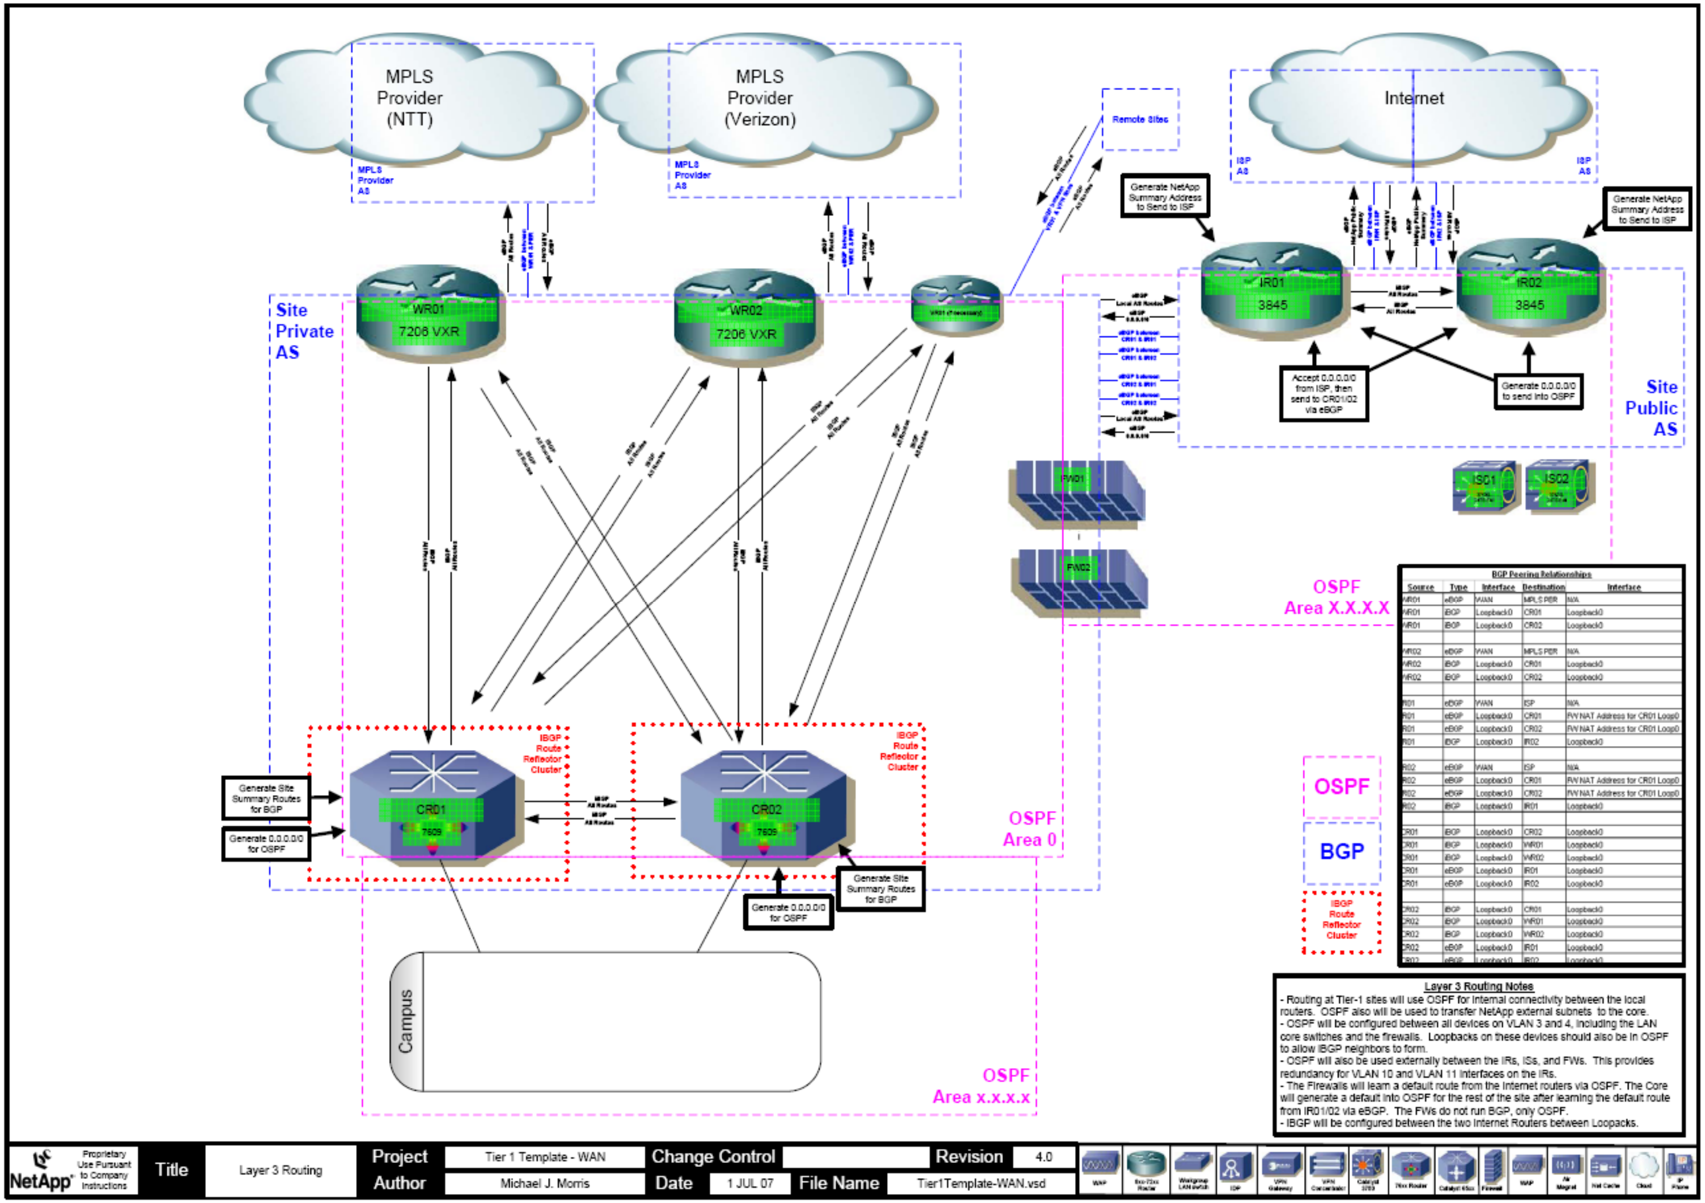
MPLS
Provider
(NTT)
MPLS Provider AS
MPLS
Provider
(Verizon)
MPLS Provider AS
Remote Sites
Internet
ISP AS
ISP AS
Generate NetApp Summary Address to Send to ISP
Generate NetApp Summary Address to Send to ISP
Site Private AS
WR01
7206 VXR
WR02
7206 VXR
IR01
3845
IR02
3845
Accept 0.0.0.0/0 from ISP, then send to CR01/02 via eBGP
Generate 0.0.0.0/0 to send into OSPF
Site Public AS
IS01
IS02
FW01
FW02
OSPF Area X.X.X.X
CR01
7609
CR02
7609
iBGP Route Reflector Cluster
iBGP Route Reflector Cluster
Generate Site Summary Routes for BGP
Generate 0.0.0.0/0 for OSPF
Generate Site Summary Routes for BGP
Generate 0.0.0.0/0 for OSPF
OSPF Area 0
Campus
OSPF Area x.x.x.x
OSPF
BGP
iBGP Route Reflector Cluster
WAP
Router
IDP
VPN Gateway
VPN Concentrator
Catalyst 3750
7600 Router
Catalyst 65xx
Firewall
WAP
Air Magnet
Net Cache
Cloud
IP Phone

BGP Peering Relationships

Source	Type	Interface	Destination	Interface
WR01	eBGP	WAN	MPLS PER	N/A
WR01	iBGP	Loopback0	CR01	Loopback0
WR01	iBGP	Loopback0	CR02	Loopback0
WR02	eBGP	WAN	MPLS PER	N/A
WR02	iBGP	Loopback0	CR01	Loopback0
WR02	iBGP	Loopback0	CR02	Loopback0
IR01	eBGP	WAN	ISP	N/A
IR01	eBGP	Loopback0	CR01	FW NAT Address for CR01 Loop0
IR01	eBGP	Loopback0	CR02	FW NAT Address for CR01 Loop0
IR01	iBGP	Loopback0	IR02	Loopback0
IR02	eBGP	WAN	ISP	N/A
IR02	eBGP	Loopback0	CR01	FW NAT Address for CR01 Loop0
IR02	eBGP	Loopback0	CR02	FW NAT Address for CR01 Loop0
IR02	iBGP	Loopback0	IR01	Loopback0
CR01	iBGP	Loopback0	CR02	Loopback0
CR01	iBGP	Loopback0	WR01	Loopback0
CR01	iBGP	Loopback0	WR02	Loopback0
CR01	eBGP	Loopback0	IR01	Loopback0
CR01	eBGP	Loopback0	IR02	Loopback0
CR02	iBGP	Loopback0	CR01	Loopback0
CR02	iBGP	Loopback0	WR01	Loopback0
CR02	iBGP	Loopback0	WR02	Loopback0
CR02	eBGP	Loopback0	IR01	Loopback0
CR02	eBGP	Loopback0	IR02	Loopback0

Layer 3 Routing Notes

- Routing at Tier-1 sites will use OSPF for internal connectivity between the local routers. OSPF also will be used to transfer NetApp external subnets to the core.
- OSPF will be configured between all devices on VLAN 3 and 4, including the LAN core switches and the firewalls. Loopbacks on these devices should also be in OSPF to allow IBGP neighbors to form.
- OSPF will also be used externally between the IRs, ISs, and FWs. This provides redundancy for VLAN 10 and VLAN 11 interfaces on the IRs.
- The Firewalls will learn a default route from the Internet routers via OSPF. The Core will generate a default into OSPF for the rest of the site after learning the default route from IR01/02 via eBGP. The FWs do not run BGP, only OSPF.
- IBGP will be configured between the two Internet Routers between Loopacks.

NetApp Proprietary Use Pursuant to Company Instructions	Title	Layer 3 Routing	Project	Tier 1 Template - WAN	Change Control		Revision	4.0
			Author	Michael J. Morris	Date	1 JUL 07	File Name	Tier1Template-WAN.vsd

NetApp™

Go further, faster

The New Global MPLS Network

Legacy WAN Issues

- To reliant on costly private lines.
- No bandwidth scalability beyond T-1 for field sites.
- Backup done by tunnels over Internet.
- Many large sites only using tunnels over Internet.
- Old hardware and no software standards.
- No QoS.
- Decentralized Internet access.
- Single Global OSPF autonomous system.
- No architectural and configuration standards.
- All sites built to different designs.

RFP Process

(9 months - start to carrier selection)

- Internal requirements and goals definition.
- Issued 2-page RFI to all possible carriers with face-to-face sales meetings with carriers performing well on RFI.
- Issued written RFP to selected carriers (AT&T, Sprint, Verizon Business, IBM, Masergy, NTT, BT).
- Carriers Rated in Three Areas.
  - Weighted Requirements Scoring Matrix.
  - Pricing and ROI.
  - Best overall value.
- Verizon Business selected as single carrier for NA and EMEA (APAC on order now).
- Benefits.
  - The power of a written RFP to refer to and rate carriers against.
  - Three completive pricing rounds led to a dual DS-3 Ethernet circuits at all NA sites.
  - Global contract centralization.
  - Leveraging a larger purchase with a single carrier (instead of a dual-carrier strategy) to drive better ROI.

The New Network

NetApp™

What We Wanted

- Global MPLS IP VPN Service
- End-to-End QoS
- WAN Ethernet Access
- Routing Protocol Scalability
- Hardware and software standards
- Sites built to approved templates
- Cost reduction
- WAN delay reductions
- More bandwidth
- Bandwidth scalability
- Diversity and SONET protection for sites
- Consolidated Internet access.

What We Got

- Verizon Private IP
- 6-Level QoS Model
- Ethernet circuit access at all NA sites and EMEA hubs
- Carrier-class BGP routing
- Complete field site HW refresh, all built to template
- Reduced telecom costs by 18% with 6-month ROI
- Cut delay with end-to-end SLAs
- Quadrupled bandwidth to NA sites.
- Able to increase MPLS port bandwidth in 7-days.
- SONET rings at all Tier-1 and Tier-2 sites.
- Centralized Internet access provided by four global hub sites with dynamic failover.

The New Global Network

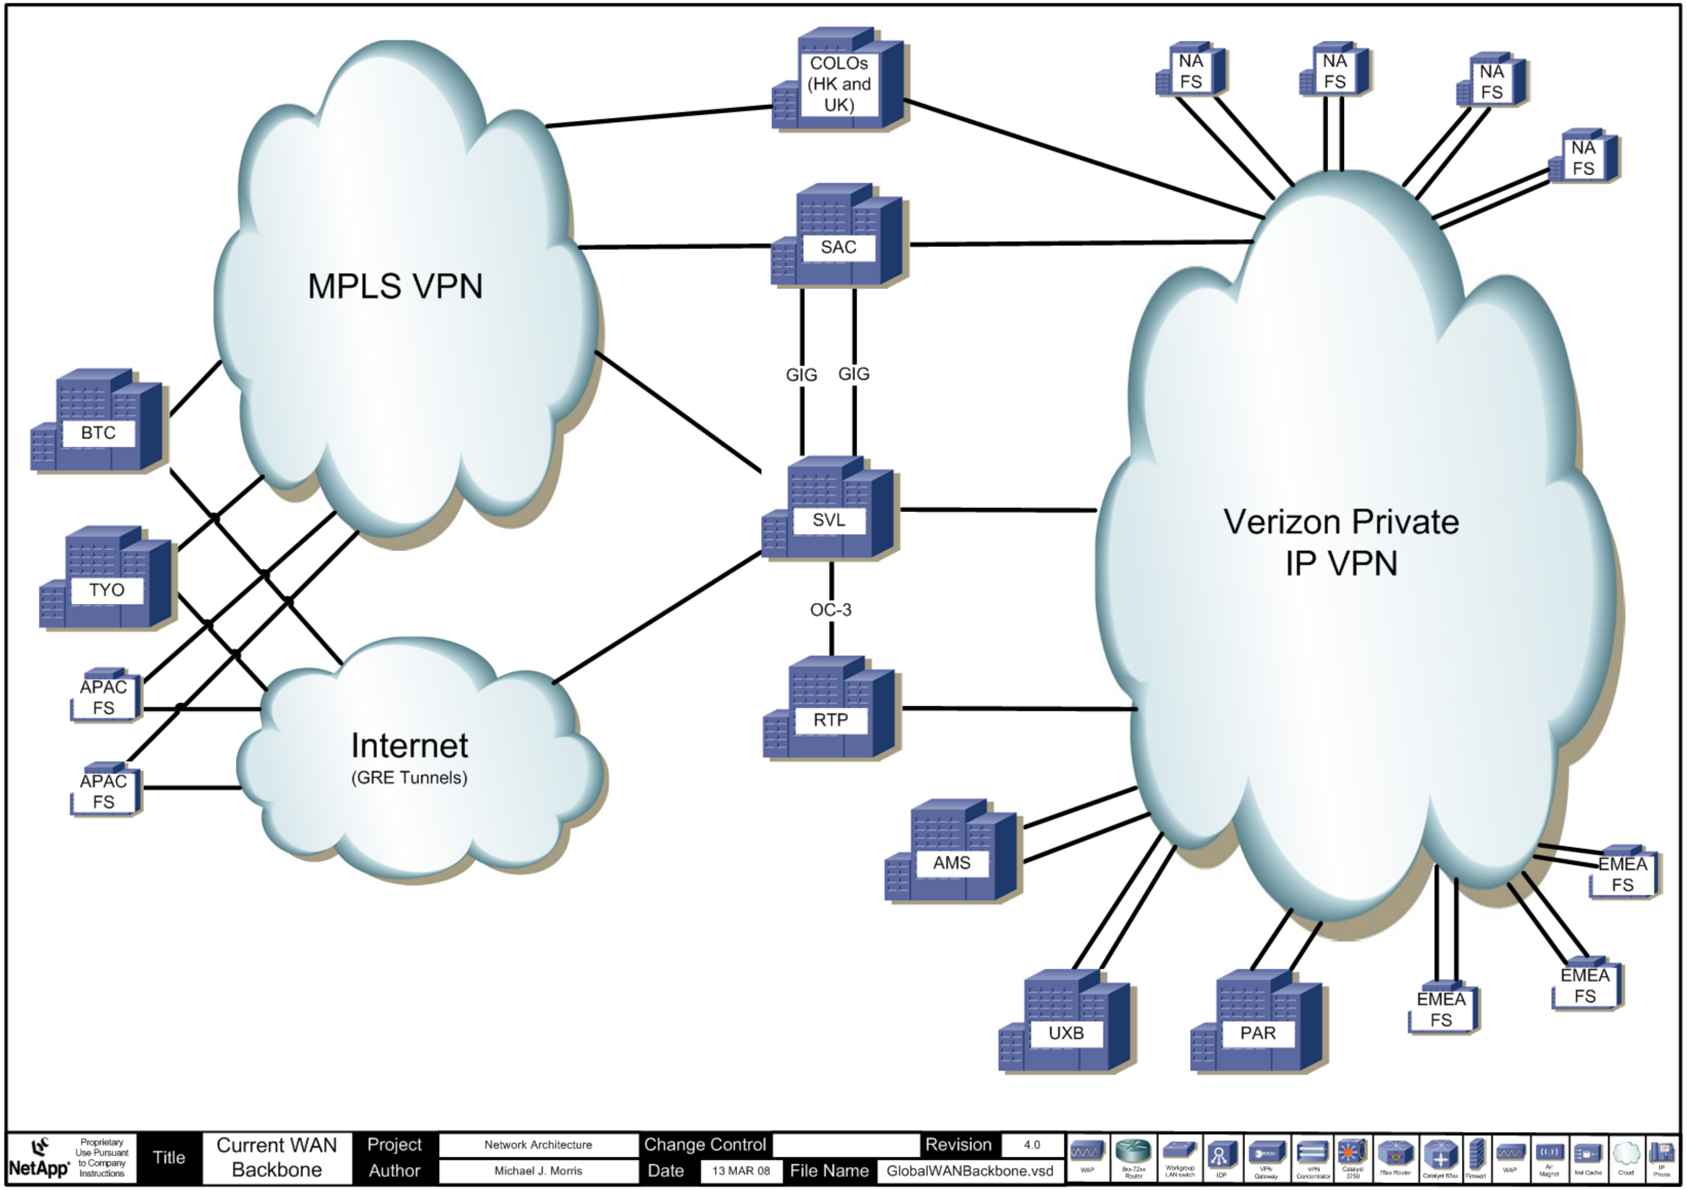

NetApp® Proprietary Use Pursuant to Company Instructions	Title	Current WAN Backbone	Project	Network Architecture	Change Control		Revision	4.0
			Author	Michael J. Morris	Date	13 MAR 08	File Name	GlobalWANBackbone.vsd

WAN Ethernet Access and DS-3s

- Most preferred local access option.
- Key feature in selecting Verizon Private IP.
- Delivered via TDM DS-3 converted to Ethernet by BEAS MUX at NetApp Office.

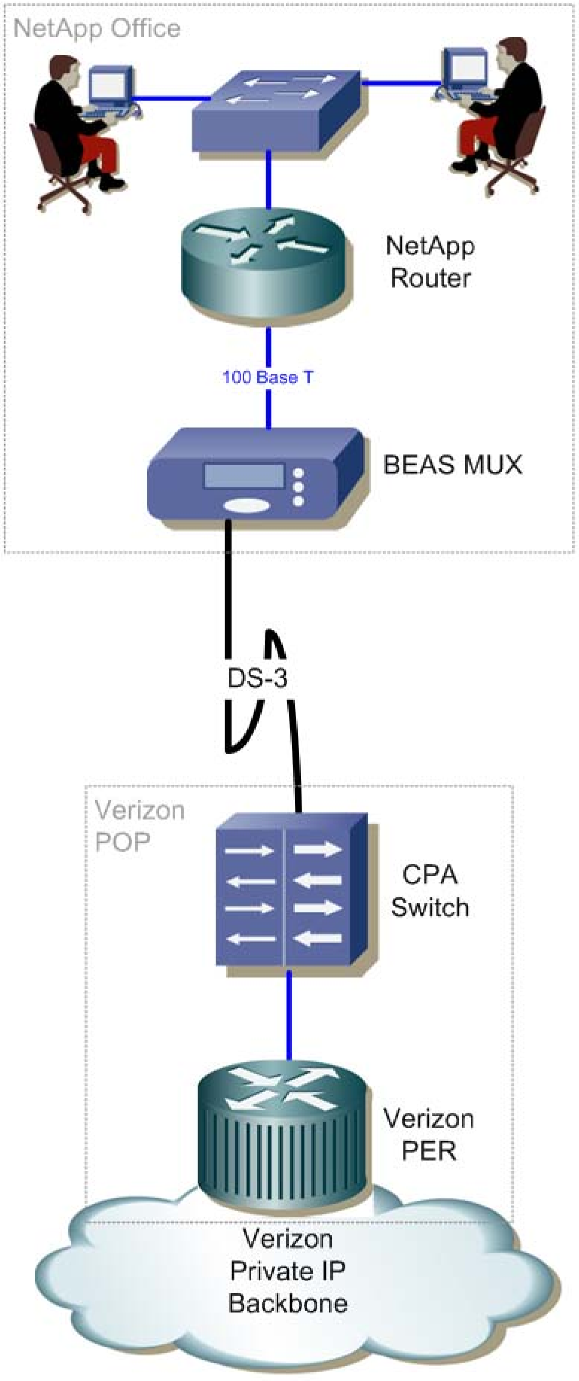

Ethernet Benefits

- Simplified connectivity.
- Bandwidth scalability.
- Lower hardware costs.
- Ubiquitous understanding by engineers.
- Virtualization (delivery of different VPNs via dot1q on the same circuit).

Routing

- BGP is now our backbone protocol.
- OSPF relegated to LAN routing and supporting iBGP establishment.
- No redistribution between routing protocols.
- Each site in its own BGP AS running eBGP with Verizon or eBGP with other sites on P2P links.
- Incredible scalability.
- Dynamic routing throughout – even to Internet routers.

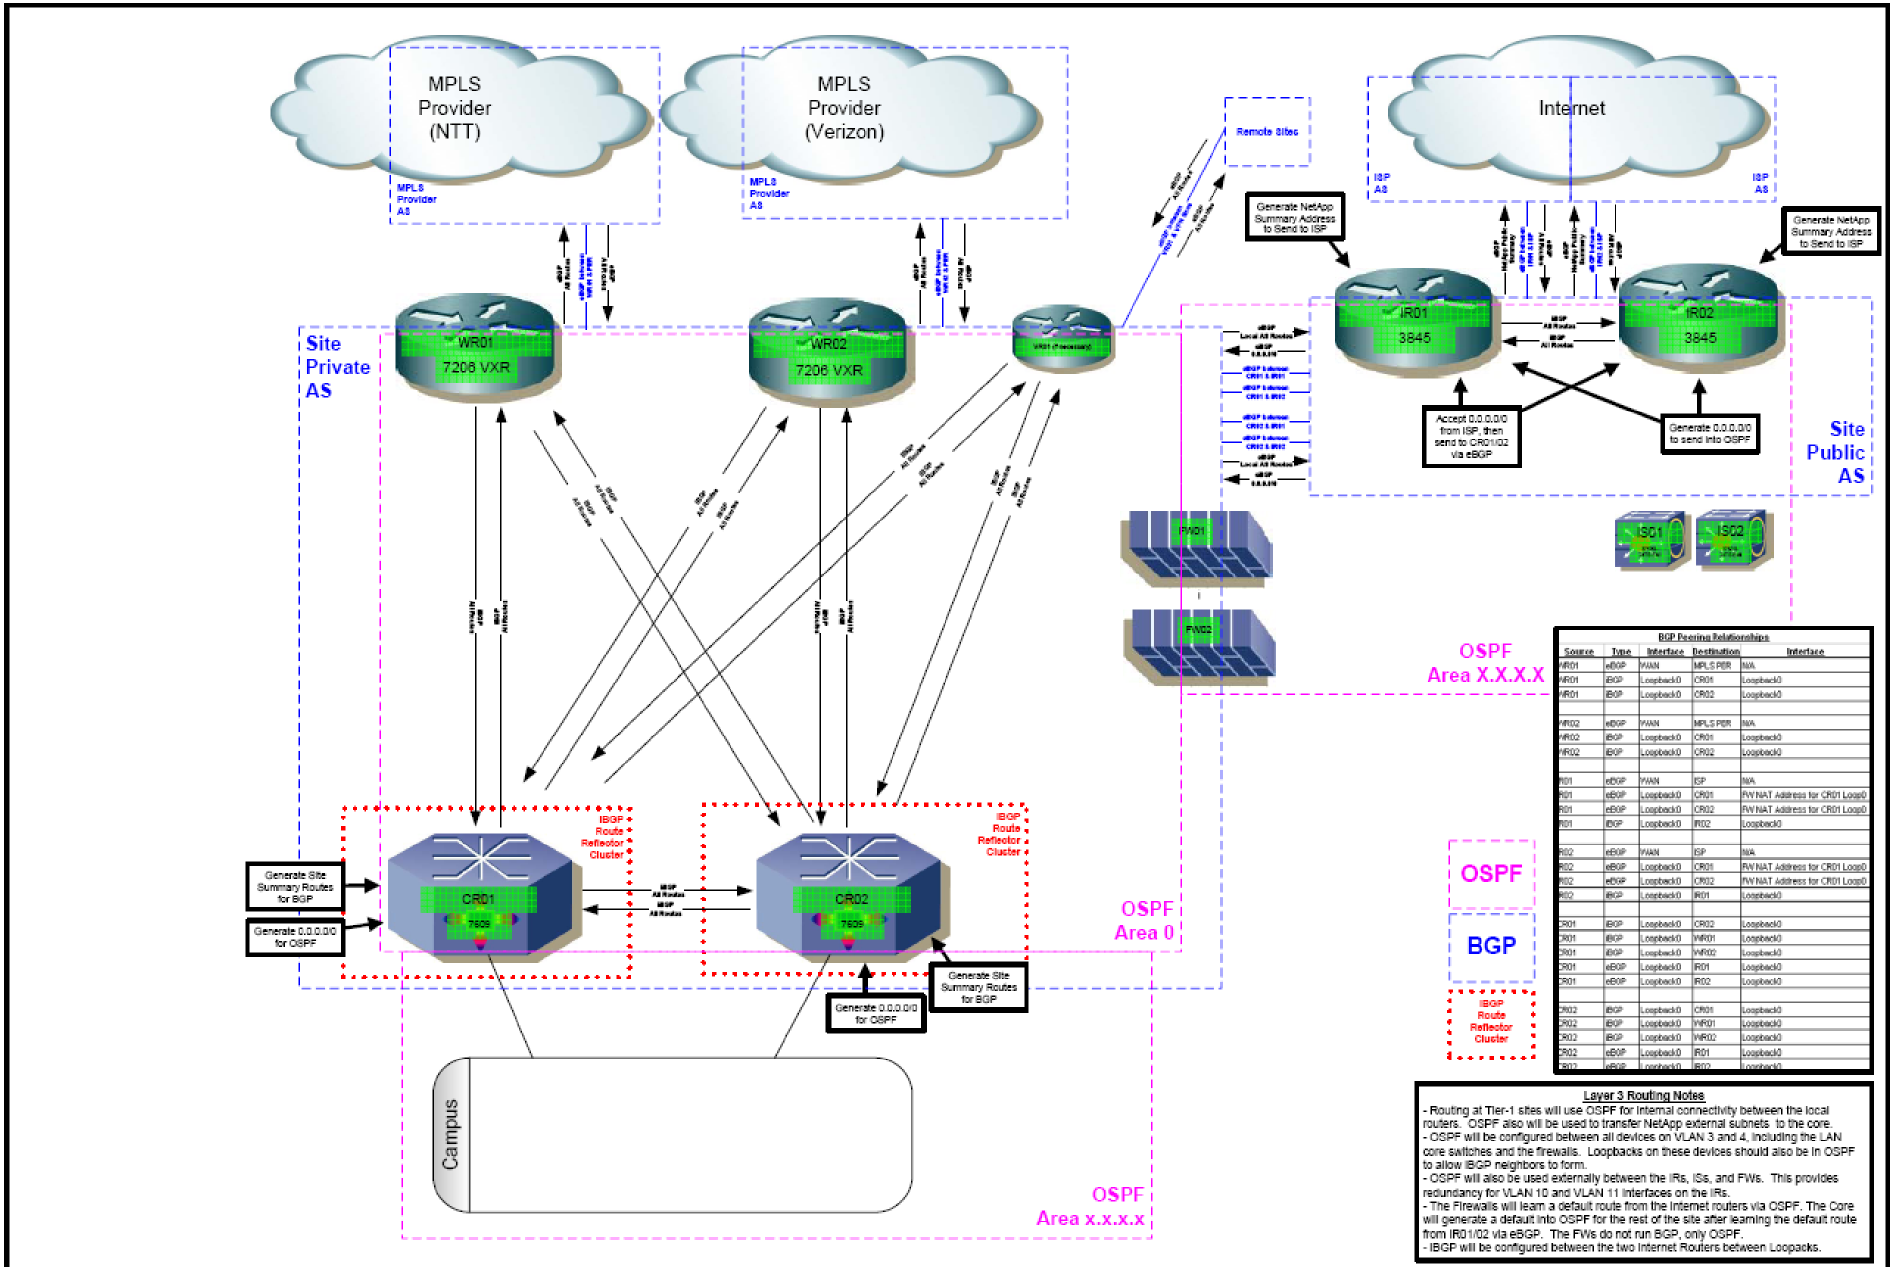
MPLS Provider (NTT)
MPLS Provider AS
MPLS Provider (Verizon)
MPLS Provider AS
Remote Sites
Internet
ISP AS
ISP AS
Generate NetApp Summary Address to Send to ISP
Generate NetApp Summary Address to Send to ISP
Site Private AS
WR01
7206 VXR
WR02
7206 VXR
VR01 (if necessary)
IR01
3845
IR02
3845
Accept 0.0.0.0/0 from ISP, then send to CR01/02 via eBGP
Generate 0.0.0.0/0 to send into OSPF
Site Public AS
FW01
FW02
IS01
IS02
OSPF Area X.X.X.X
eBGP All Routes
iBGP All Routes
IBGP Route Reflector Cluster
IBGP Route Reflector Cluster
CR01
7609
CR02
7609
Generate Site Summary Routes for BGP
Generate 0.0.0.0/0 for OSPF
Generate 0.0.0.0/0 for OSPF
Generate Site Summary Routes for BGP
OSPF Area 0
Campus
OSPF Area x.x.x.x
OSPF
BGP
IBGP Route Reflector Cluster

BGP Peering Relationships

Source	Type	Interface	Destination	Interface
WR01	eBGP	WAN	MPLS PER	N/A
WR01	iBGP	Loopback0	CR01	Loopback0
WR01	iBGP	Loopback0	CR02	Loopback0
WR02	eBGP	WAN	MPLS PER	N/A
WR02	iBGP	Loopback0	CR01	Loopback0
WR02	iBGP	Loopback0	CR02	Loopback0
IR01	eBGP	WAN	ISP	N/A
IR01	eBGP	Loopback0	CR01	FW NAT Address for CR01 Loop0
IR01	eBGP	Loopback0	CR02	FW NAT Address for CR01 Loop0
IR01	iBGP	Loopback0	IR02	Loopback0
IR02	eBGP	WAN	ISP	N/A
IR02	eBGP	Loopback0	CR01	FW NAT Address for CR01 Loop0
IR02	eBGP	Loopback0	CR02	FW NAT Address for CR01 Loop0
IR02	iBGP	Loopback0	IR01	Loopback0
CR01	iBGP	Loopback0	CR02	Loopback0
CR01	iBGP	Loopback0	WR01	Loopback0
CR01	iBGP	Loopback0	WR02	Loopback0
CR01	eBGP	Loopback0	IR01	Loopback0
CR01	eBGP	Loopback0	IR02	Loopback0
CR02	iBGP	Loopback0	CR01	Loopback0
CR02	iBGP	Loopback0	WR01	Loopback0
CR02	iBGP	Loopback0	WR02	Loopback0
CR02	eBGP	Loopback0	IR01	Loopback0
CR02	eBGP	Loopback0	IR02	Loopback0

Layer 3 Routing Notes

- Routing at Tier-1 sites will use OSPF for internal connectivity between the local routers. OSPF also will be used to transfer NetApp external subnets to the core.
- OSPF will be configured between all devices on VLAN 3 and 4, including the LAN core switches and the firewalls. Loopbacks on these devices should also be in OSPF to allow IBGP neighbors to form.
- OSPF will also be used externally between the IRs, ISs, and FWs. This provides redundancy for VLAN 10 and VLAN 11 interfaces on the IRs.
- The Firewalls will learn a default route from the Internet routers via OSPF. The Core will generate a default into OSPF for the rest of the site after learning the default route from IR01/02 via eBGP. The FWs do not run BGP, only OSPF.
- IBGP will be configured between the two Internet Routers between Loopacks.

NetApp® Proprietary Use Pursuant to Company Instructions

Title	Layer 3 Routing	Project	Tier 1 Template - WAN	Change Control	
				Revision	4.0
		Author	Michael J. Morris	Date	1 JUL 07
				File Name	Tier1Template-WAN.vsd

Dynamic Internet Access

- BGP peering between internal core routers and NetApp Internet routers.
- Allows hub sites and large offices to dynamically accept default route from ISPs and pass default to internal network.
  - Internal core then updates Verizon Private IP via BGP.
  - Field sites route dynamically over Verizon Private IP to nearest hub site for Internet access.
- If a hub or large site’s Internet circuit goes down:
  - BGP updates internal core, which updates Verizon Private IP, which reconverges to another hub site.
  - Reconvergence takes about 8 seconds.
- Internet circuits at large sites reduced to one circuit, field sites use hub sites.
- Reduced ISP costs by 61%.

Lessoned Learned

- Standards and templates come first. Know what you want to build before trying to build it.
- A written RFP provides so many benefits. Take the time to do one.
- Build it as a global team. It takes a lot to do a global network.
- BGP is the ***new*** enterprise protocol. Use it and all of its features.
- Make everything dynamic – a good network design can make all parts of your network dynamic and reduce your costs.
- Aggressive pricing and carrier credits can pay for legacy circuit cancellations.
- Specify required carrier diversity at all levels (circuit, CO, long-distance backhaul, POP, PER). Then, be completely involved in the carrier's circuit provisioning.
- Insist on end-to-end, per-site SLAs. Carriers, by default, only provide averaged backbone SLAs. Those are useless.
- The cheapest is not the best. Evaluate on more than cost.

NetApp™

Go further, faster

The Future

Future Strategies and Needs

- Dynamic changes with providers - QoS, BW, etc.
- Faster provisioning.
- Looking glasses with MPLS providers.
- Routing Virtualization.
- Application acceleration.
- Extranets.
- More Ethernet access.
- MPLS "Internet“ with Standard QoS.

Questions ?

NetApp™

Go further, faster

NetApp Labs MPLS IP VPN Network (if time available)

Reasons for Building

- NetApp software developers place huge demands on corporate network for bandwidth.
- Several outages caused by rouge lab traffic.
- Isolate lab traffic to protect corporate assets.
- Provide optimal connectivity between technology center labs.

High Level Design

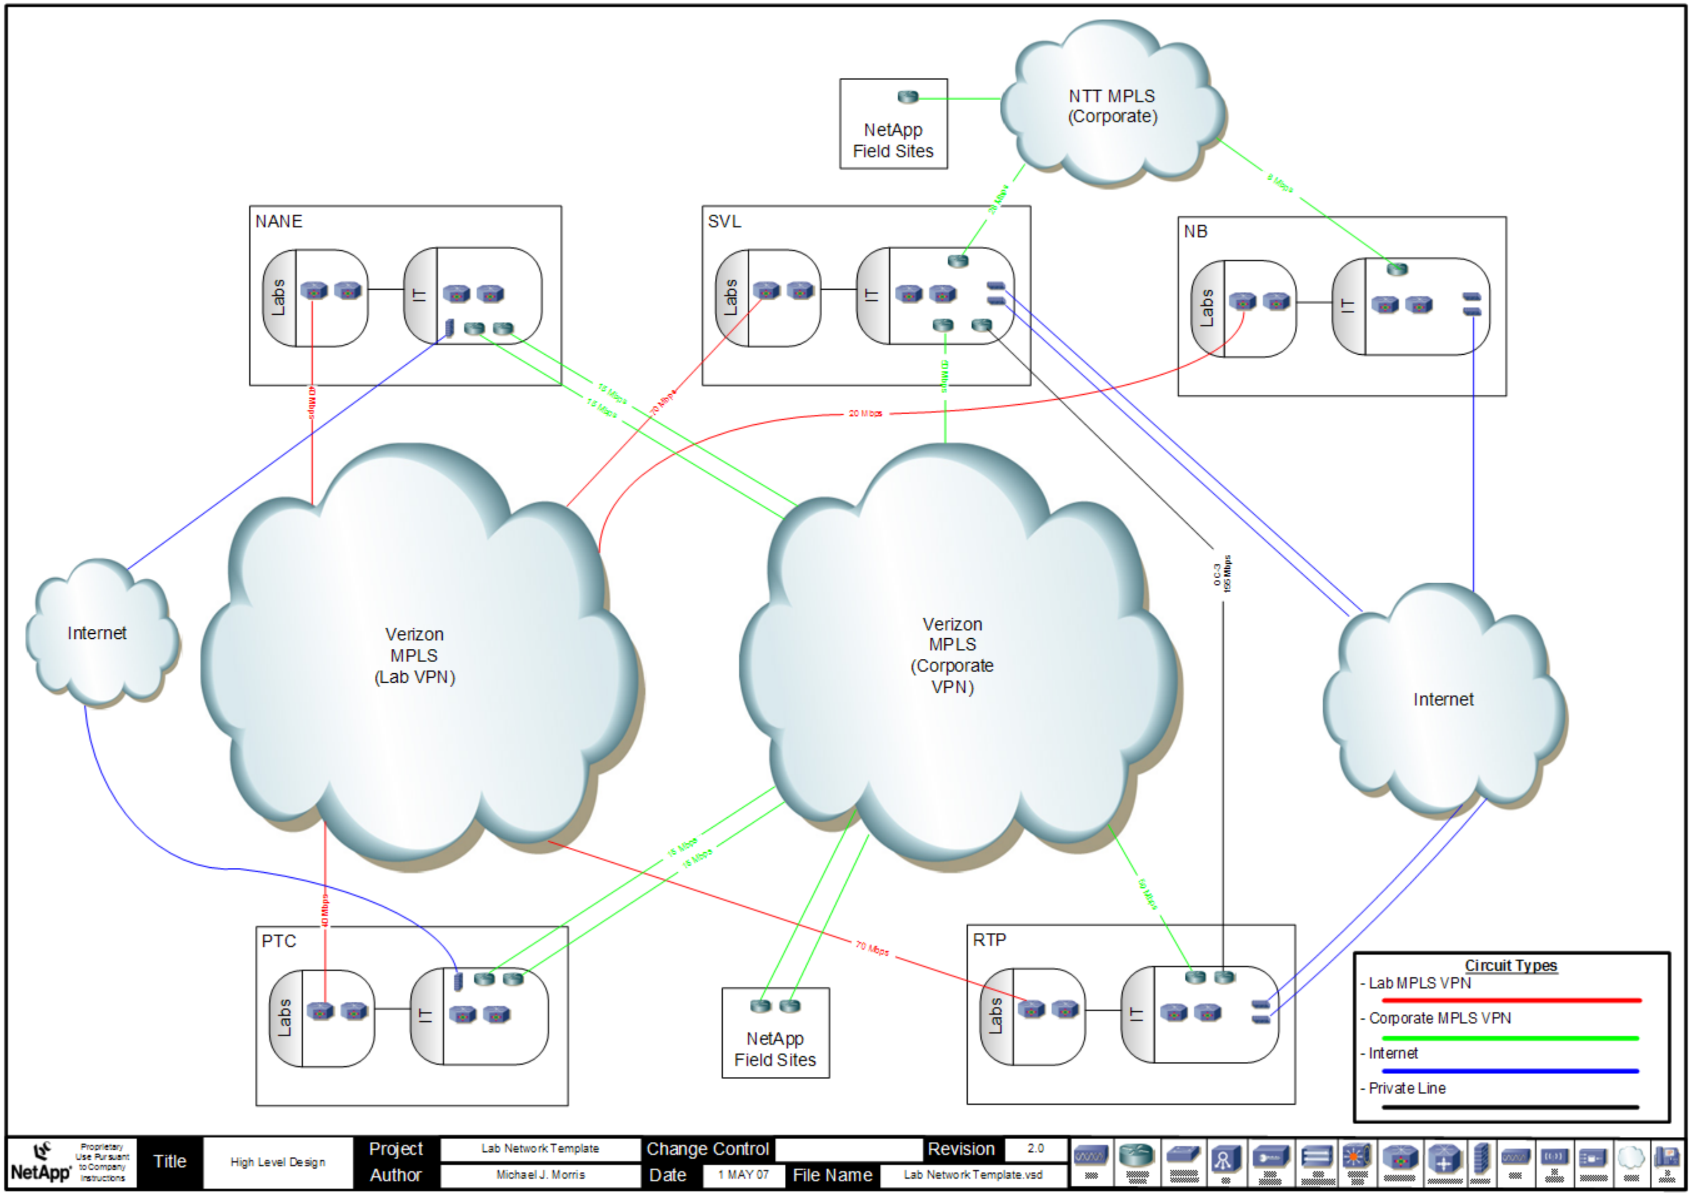

Title	High Level Design	Project	Lab Network Template	Change Control		Revision	2.0
		Author	Michael J. Morris	Date	1 MAY 07	File Name	Lab Network Template.vsd

Lab Routing

- All based on BGP as part of corporate backbone routing protocol.
- All labs consolidated behind IT Lab Core routers which are in their own BGP AS.
  - eBGP to IT core.
  - eBGP to Verizon MPLS (Lab VRF).
  - eBGP or OSPF to the labs.
  - Labs inside unique set of BGP AS's to identify routes.
- After identifying routes generated by labs using BGP AS number, BGP controls used to influence routing.
- Result – as long as one end of a session is in a lab, even when the other end of the session in in the IT network, all traffic flows over the Lab MPLS network.
  - No policy routing required, all done with BGP.
  - Single large links for Lab MPLS with backup via Corporate MPLS.

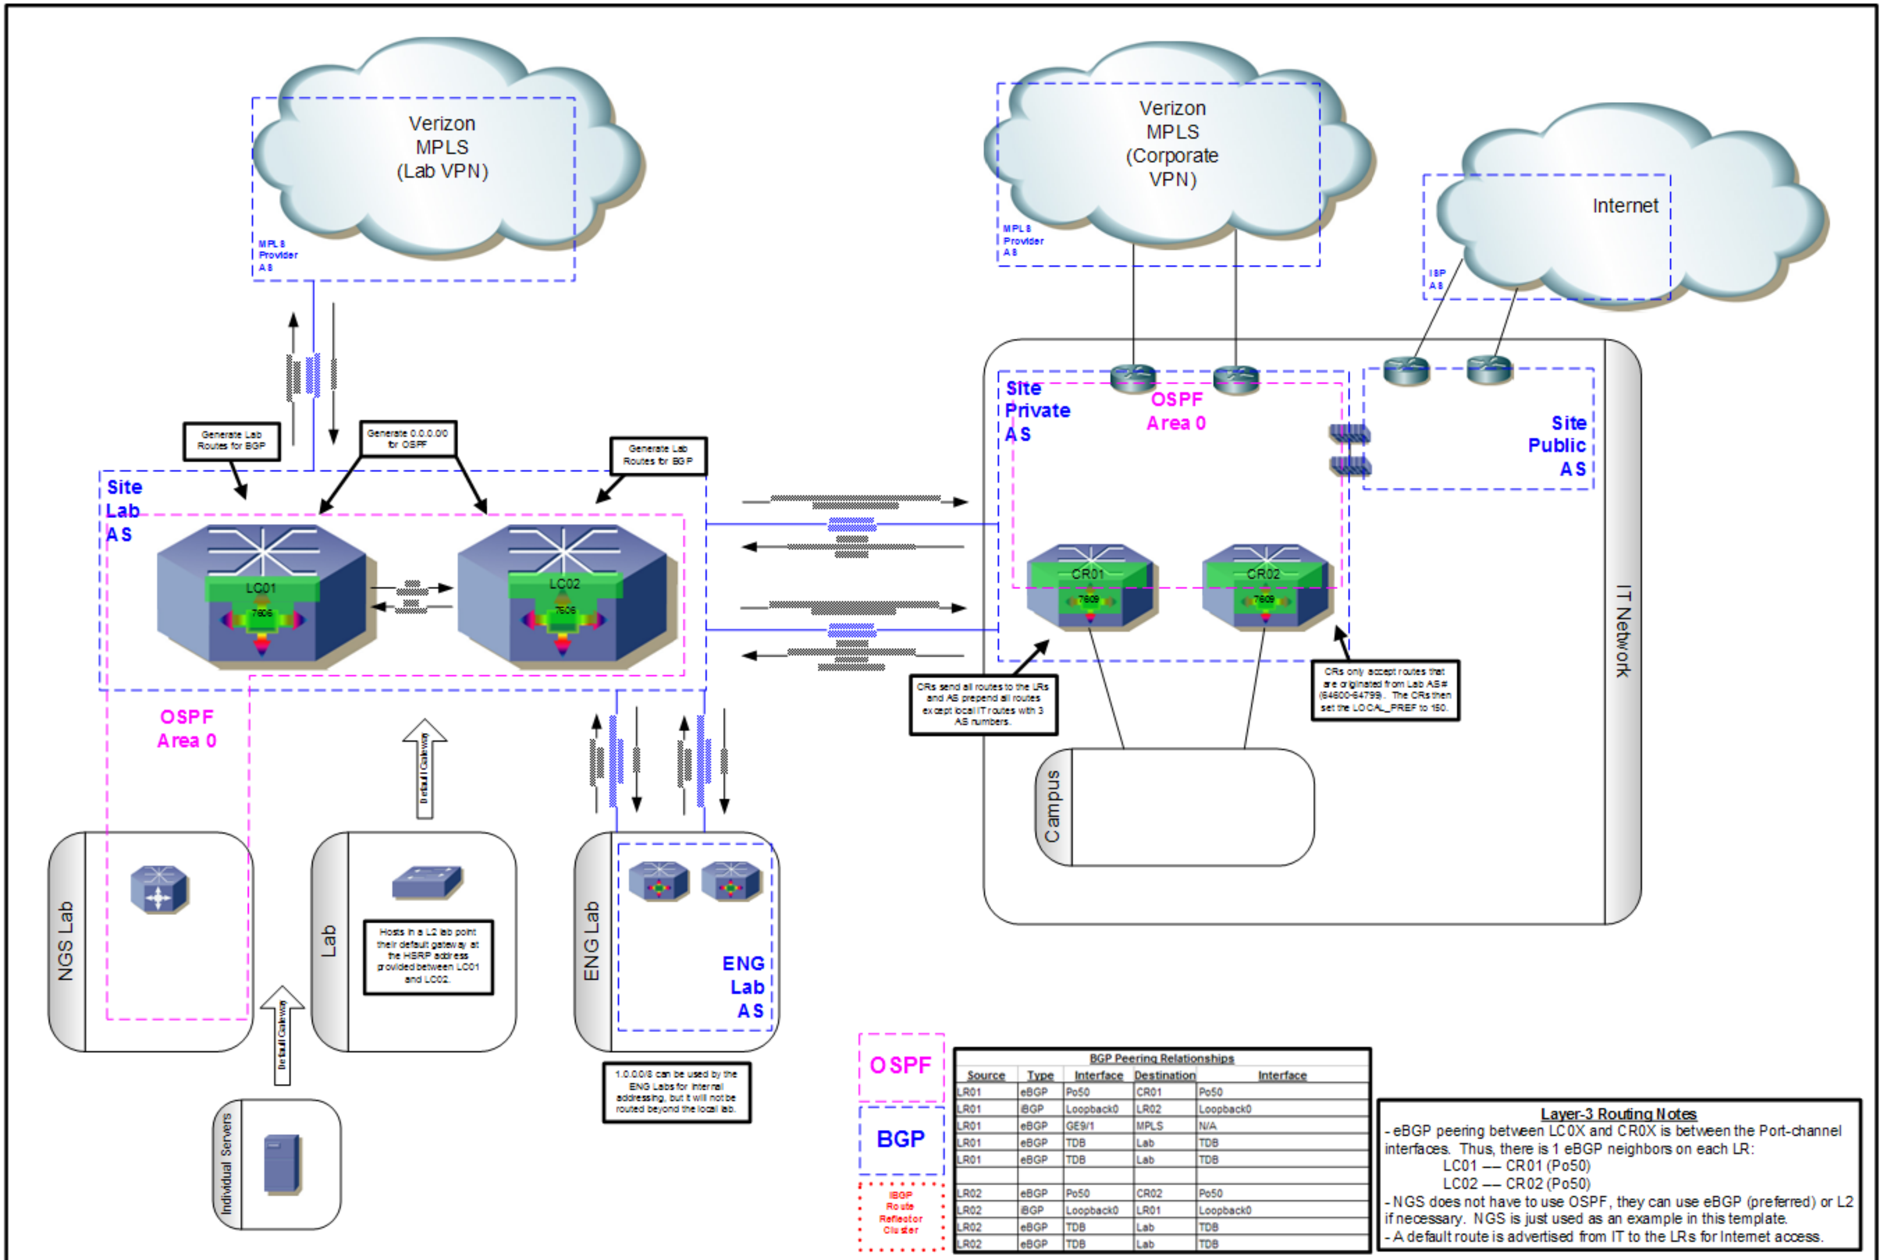

Generate Lab Routes for BGP

Generate 0.0.0.0/0 for OSPF

Generate Lab Routes for BGP

CRs send all routes to the LRs and AS prepend all routes except local IT routes with 3 AS numbers.

CRs only accept routes that are originated from Lab AS# (64600-64799). The CRs then set the LOCAL_PREF to 150.

Hosts in a L2 lab point their default gateway at the HSRP address provided between LC01 and LC02.

1.0.0.0/8 can be used by the ENG Labs for internal addressing, but it will not be routed beyond the local lab.

BGP Peering Relationships

Source	Type	Interface	Destination	Interface
LR01	eBGP	Po50	CR01	Po50
LR01	iBGP	Loopback0	LR02	Loopback0
LR01	eBGP	GE9/1	MPLS	N/A
LR01	eBGP	TDB	Lab	TDB
LR01	eBGP	TDB	Lab	TDB
LR02	eBGP	Po50	CR02	Po50
LR02	iBGP	Loopback0	LR01	Loopback0
LR02	eBGP	TDB	Lab	TDB
LR02	eBGP	TDB	Lab	TDB

Layer-3 Routing Notes

- eBGP peering between LC0X and CR0X is between the Port-channel interfaces. Thus, there is 1 eBGP neighbors on each LR:
  LC01 –– CR01 (Po50)
  LC02 –– CR02 (Po50)
- NGS does not have to use OSPF, they can use eBGP (preferred) or L2 if necessary. NGS is just used as an example in this template.
- A default route is advertised from IT to the LRs for Internet access.

NetApp Proprietary Use Pursuant to Company Instructions	Title	Layer-3 Routing	Project	Lab Network Template	Change Control		Revision	2.0
			Author	Michael J. Morris	Date	1 MAY 07	File Name	Lab Network Template.vsd

Lab Traffic Flows

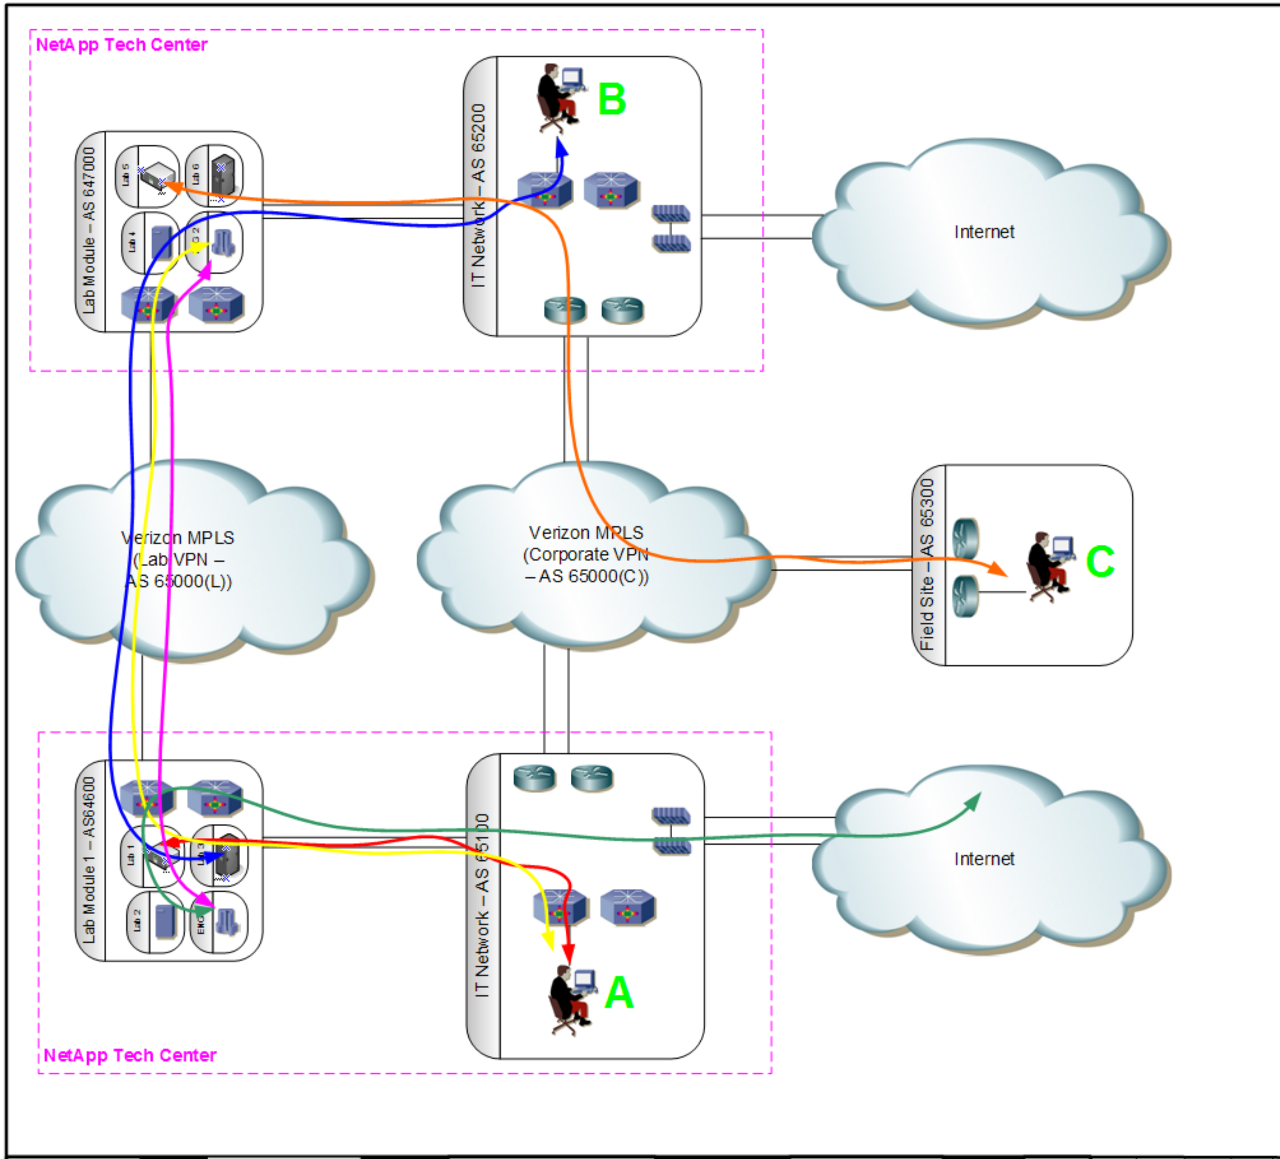

Traffic Flows

- User "A" to/from "Lab 1" ⟷
- User "A" to/from "ENG 2" ⟷
- User "B" to/from "Lab 3" ⟷
- "ENG 1" to/from Internet ⟷
- "ENG 1" to/from "ENG 2" ⟷
- User "C" to/from "Lab 5" ⟷

Traffic Flow Requirements

- Lab-to-Lab traffic stays on the Lab VPN.

- Local users only use the LAN to access the local labs. These is no backup in the event all 4 local LAN links fail.

- Remote users (from their desktops in the IT network) access remote labs over the Lab MPLS VPN. Return traffic flows over the Lab VPN.

- The Lab VPN does not provide backup for the Corporate MPLS VPN.

NetApp Proprietary Use Pursuant to Company Instructions	Title	Traffic Flows	Project	Lab Network Template	Change Control		Revision 2.0
			Author	Michael J. Morris	Date	1 MAY 07	File Name Lab Network Template.vsd

Backup Lab Traffic Flows

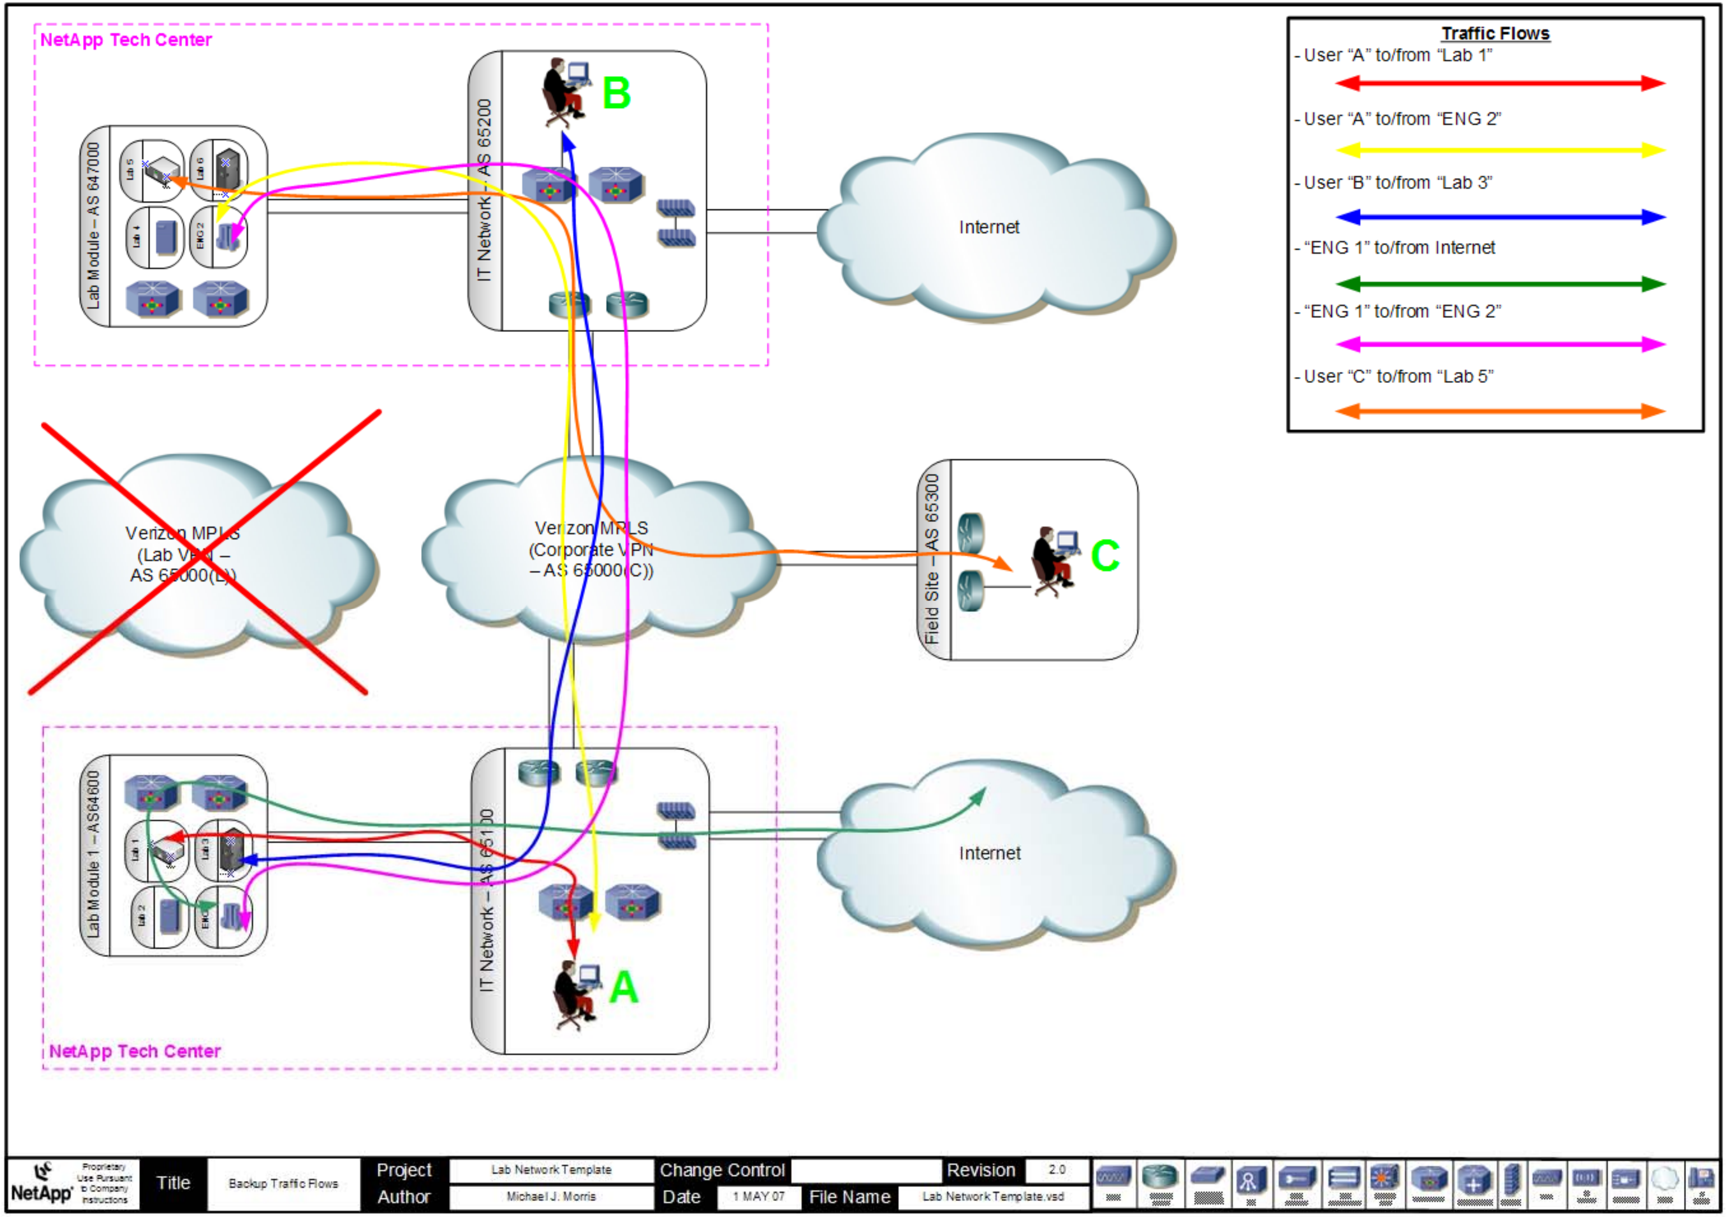

Traffic Flows

- User "A" to/from "Lab 1" 🟥
- User "A" to/from "ENG 2" 🟨
- User "B" to/from "Lab 3" 🟦
- "ENG 1" to/from Internet 🟩
- "ENG 1" to/from "ENG 2" 🟪
- User "C" to/from "Lab 5" 🟧

NetApp Proprietary Use Pursuant to Company Instructions	Title	Backup Traffic Flows	Project	Lab Network Template	Change Control		Revision	2.0
			Author	Michael J. Morris	Date	1 MAY 07	File Name	Lab Network Template.vsd

Does it Work?

2

Primary Lab Router with Lab MPLS circuit goes down.

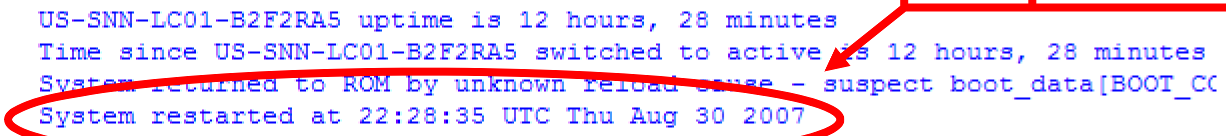

```
US-SNN-LC01-B2F2RA5 uptime is 12 hours, 28 minutes
Time since US-SNN-LC01-B2F2RA5 switched to active is 12 hours, 28 minutes
System returned to ROM by unknown reload cause - suspect boot_data[BOOT_CC
System restarted at 22:28:35 UTC Thu Aug 30 2007
```

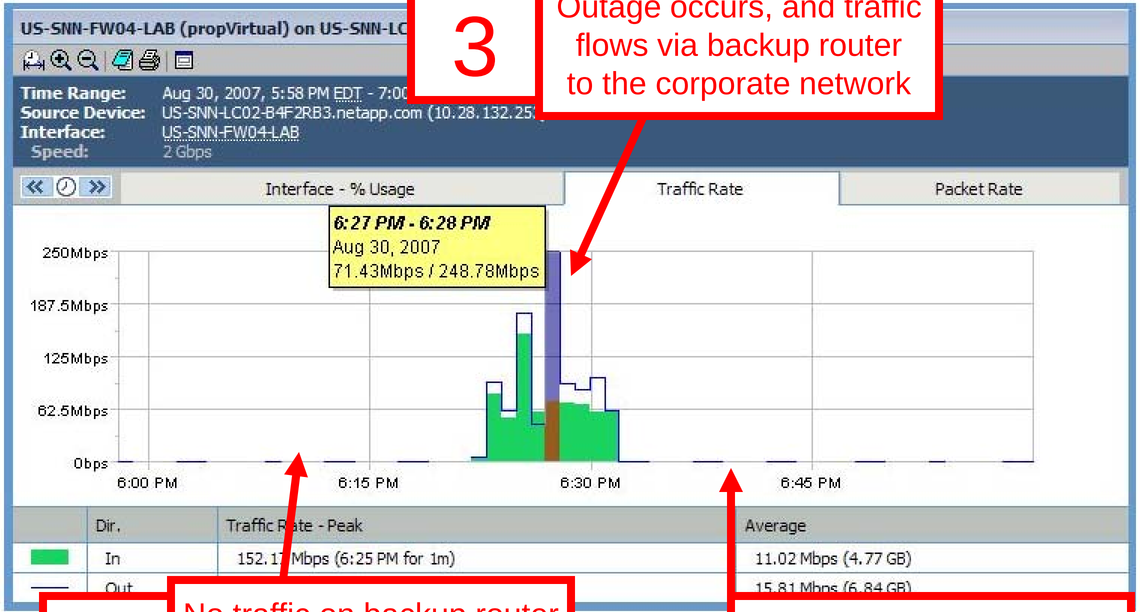

US-SNN-FW04-LAB (propVirtual) on US-SNN-LC

Time Range: Aug 30, 2007, 5:58 PM EDT - 7:00
Source Device: US-SNN-LC02-B4F2RB3.netapp.com (10.28.132.25
Interface: US-SNN-FW04-LAB
Speed: 2 Gbps

	Dir.	Traffic Rate - Peak	Average
🟩	In	152.17 Mbps (6:25 PM for 1m)	11.02 Mbps (4.77 GB)
—	Out		15.81 Mbps (6.84 GB)

3

Outage occurs, and traffic flows via backup router to the corporate network

5

During the outage, traffic between labs can be seen on corporate WAN.

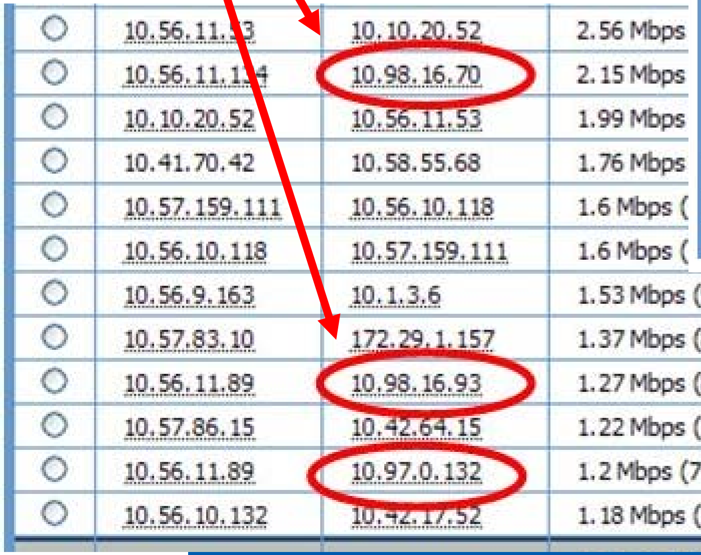

○	10.56.11.53	10.10.20.52	2.56 Mbps
○	10.56.11.134	10.98.16.70	2.15 Mbps
○	10.10.20.52	10.56.11.53	1.99 Mbps
○	10.41.70.42	10.58.55.68	1.76 Mbps
○	10.57.159.111	10.56.10.118	1.6 Mbps (
○	10.56.10.118	10.57.159.111	1.6 Mbps (
○	10.56.9.163	10.1.3.6	1.53 Mbps (98.
○	10.57.83.10	172.29.1.157	1.37 Mbps (88.
○	10.56.11.89	10.98.16.93	1.27 Mbps (81.47 MB)
○	10.57.86.15	10.42.64.15	1.22 Mbps (78.36 MB)
○	10.56.11.89	10.97.0.132	1.2 Mbps (77.26 MB)
○	10.56.10.132	10.42.17.52	1.18 Mbps (76.25 MB)

1

No traffic on backup router with links to corporate network, before outage

4

Primary router and Lab MPLS circuit restore and traffic returns to normal path.

Benefits

- Lab traffic optimized.
  - Direct WAN connectivity between labs.
  - Separate QoS policy for lab traffic.
  - Reduced delay.
- Large bandwidth links for software developers to use for testing.
- Lab traffic removed from corporate network.
- No outages caused by rogue lab traffic.
- Consolidation of lab resources into new data center in RTP, NC.
- Bangalore, India software developers much more productive.

NetApp™

Go further, faster

Extra Slides

Stuff About Me

- Worked in networking for over 10 years with a focus on global, enterprise WANs.
- Four years in the US Army as a Communications Officer and paratrooper with the 82nd Airborne.
- CCIE #11733.
- Worked on two Fortune 10 networks, government networks, and other enterprise networks.
- Designed and built several global MPLS networks.
- Featured blogger on Network World’s Cisco Subnet
  - http://www.networkworld.com/community/morris
- Serve on Cisco advisor board and customer councils.

Network Architecture (wiki)

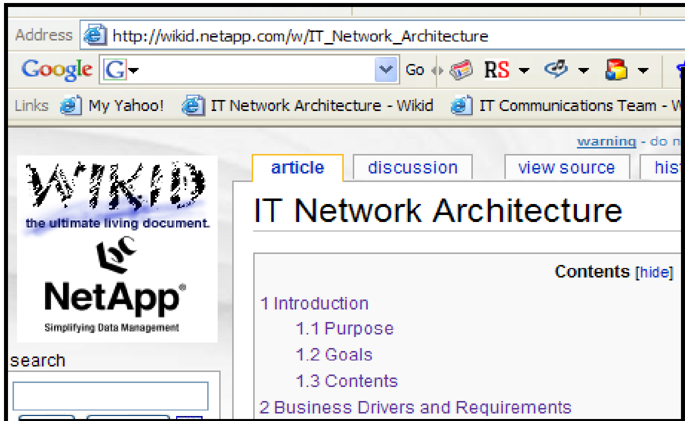

http://wikid.netapp.com/w/IT_Network_Architecture

My Yahoo! | IT Network Architecture - Wikid | IT Communications Team - W

WIKID the ultimate living document. NetApp Simplifying Data Management

IT Network Architecture

Contents [hide]

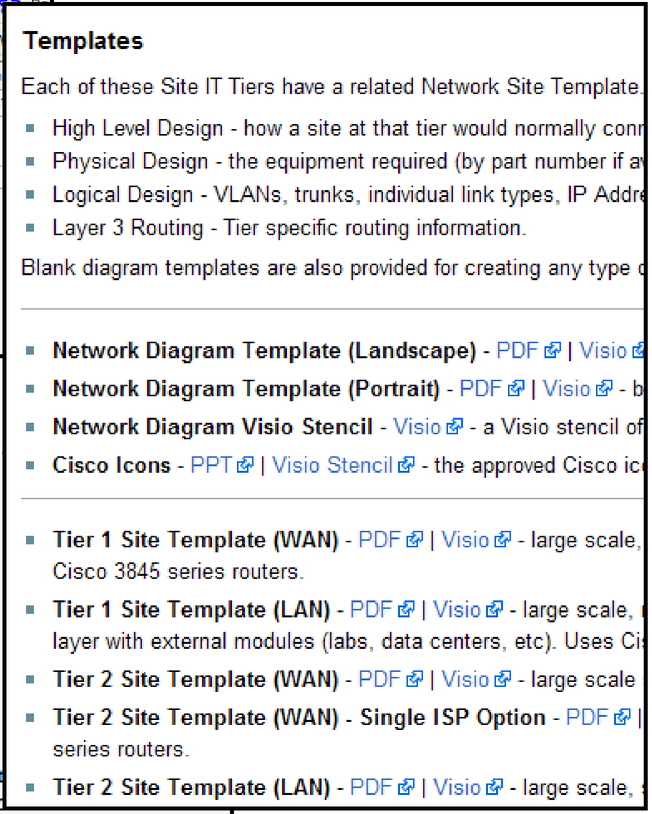

Templates

Each of these Site IT Tiers have a related Network Site Template.

- High Level Design - how a site at that tier would normally conn
- Physical Design - the equipment required (by part number if av
- Logical Design - VLANs, trunks, individual link types, IP Addre
- Layer 3 Routing - Tier specific routing information.

Blank diagram templates are also provided for creating any type o

- **Network Diagram Template (Landscape)** - PDF | Visio
- **Network Diagram Template (Portrait)** - PDF | Visio - b
- **Network Diagram Visio Stencil** - Visio - a Visio stencil of
- **Cisco Icons** - PPT | Visio Stencil - the approved Cisco ic

- **Tier 1 Site Template (WAN)** - PDF | Visio - large scale, Cisco 3845 series routers.
- **Tier 1 Site Template (LAN)** - PDF | Visio - large scale, layer with external modules (labs, data centers, etc). Uses Ci
- **Tier 2 Site Template (WAN)** - PDF | Visio - large scale
- **Tier 2 Site Template (WAN) - Single ISP Option** - PDF | series routers.
- **Tier 2 Site Template (LAN)** - PDF | Visio - large scale,

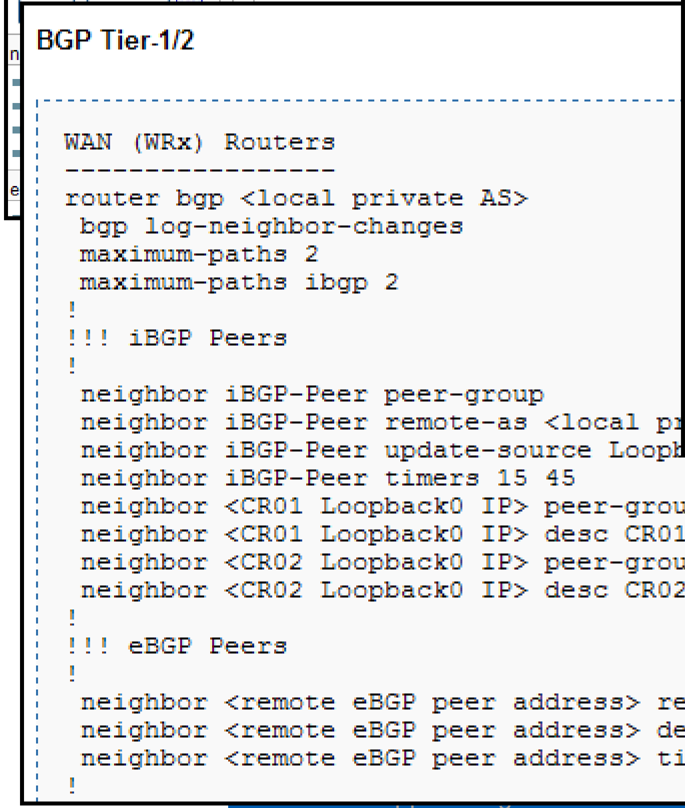

BGP Tier-1/2

```
WAN (WRx) Routers
------------------
router bgp <local private AS>
 bgp log-neighbor-changes
 maximum-paths 2
 maximum-paths ibgp 2
!
!!! iBGP Peers
!
 neighbor iBGP-Peer peer-group
 neighbor iBGP-Peer remote-as <local pr
 neighbor iBGP-Peer update-source Looph
 neighbor iBGP-Peer timers 15 45
 neighbor <CR01 Loopback0 IP> peer-group iBGP-Peer
 neighbor <CR01 Loopback0 IP> desc CR01
 neighbor <CR02 Loopback0 IP> peer-group iBGP-Peer
 neighbor <CR02 Loopback0 IP> desc CR02
!
!!! eBGP Peers
!
 neighbor <remote eBGP peer address> remote-as <remot
 neighbor <remote eBGP peer address> description <MPL
 neighbor <remote eBGP peer address> timers 15 45
!
```

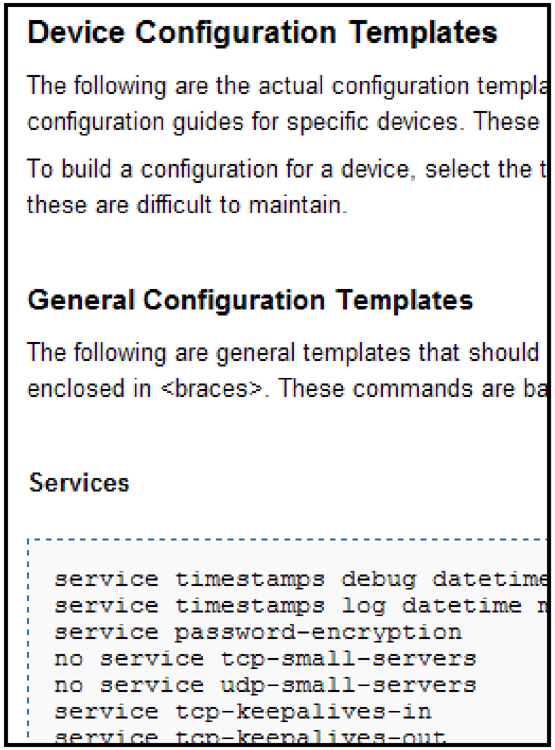

Device Configuration Templates

The following are the actual configuration templa configuration guides for specific devices. These

To build a configuration for a device, select the t these are difficult to maintain.

General Configuration Templates

The following are general templates that should enclosed in <braces>. These commands are ba

Services

```
service timestamps debug datetime
service timestamps log datetime m
service password-encryption
no service tcp-small-servers
no service udp-small-servers
service tcp-keepalives-in
service tcp-keepalives-out
```

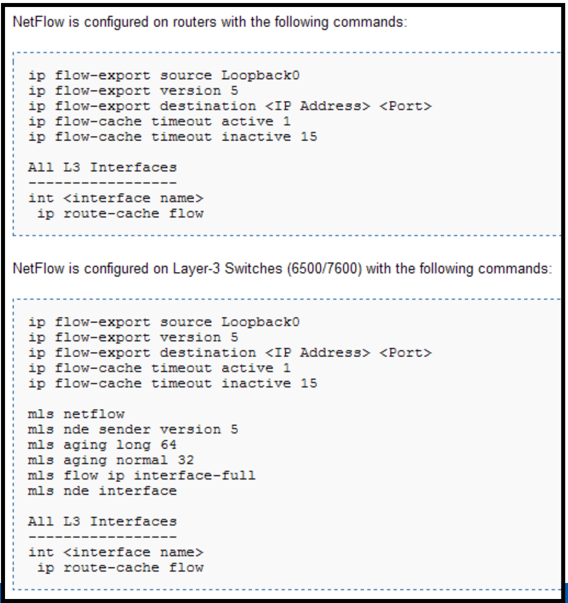

NetFlow is configured on routers with the following commands:

```
ip flow-export source Loopback0
ip flow-export version 5
ip flow-export destination <IP Address> <Port>
ip flow-cache timeout active 1
ip flow-cache timeout inactive 15

All L3 Interfaces
-----------------
int <interface name>
 ip route-cache flow
```

NetFlow is configured on Layer-3 Switches (6500/7600) with the following commands:

```
ip flow-export source Loopback0
ip flow-export version 5
ip flow-export destination <IP Address> <Port>
ip flow-cache timeout active 1
ip flow-cache timeout inactive 15

mls netflow
mls nde sender version 5
mls aging long 64
mls aging normal 32
mls flow ip interface-full
mls nde interface

All L3 Interfaces
-----------------
int <interface name>
 ip route-cache flow
```

Legacy WAN

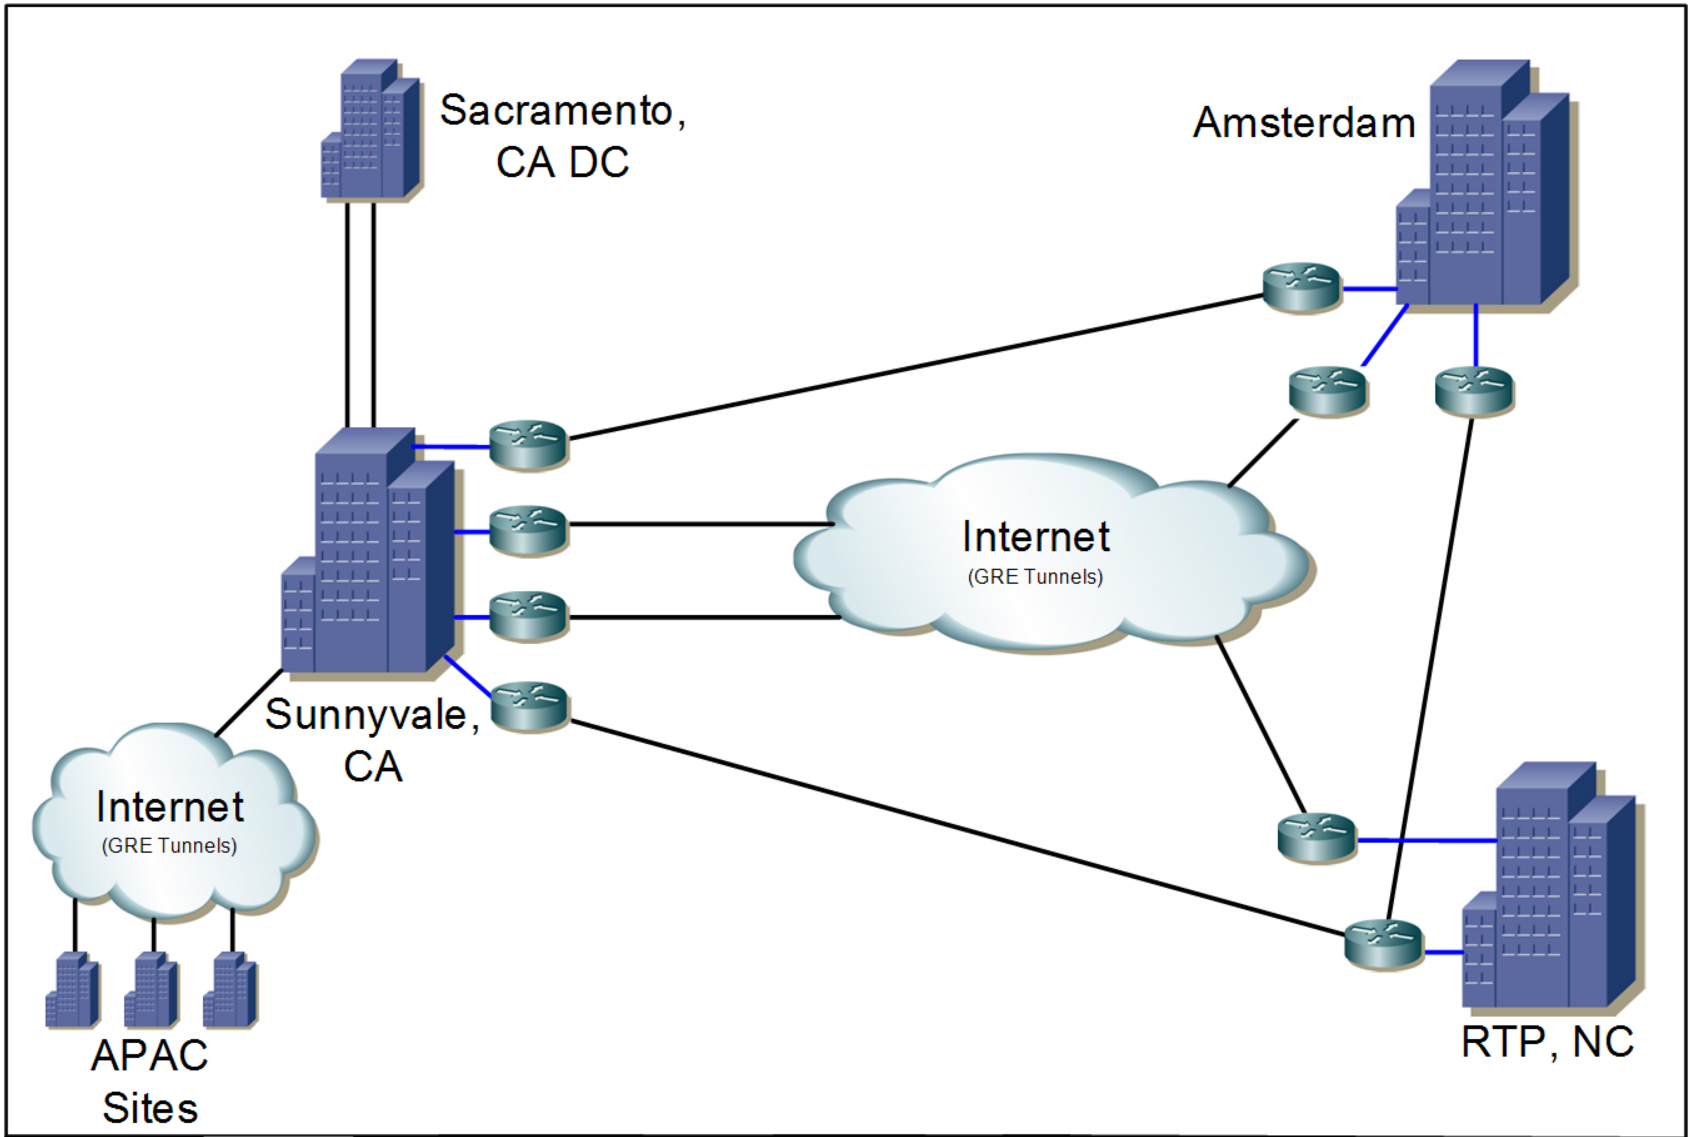

NetApp Proprietary Use Pursuant to Company Instructions	Title	Legacy Global WAN Backbone	Project	Network Architecture	Change Control		Revision	3.0
			Author	Michael J. Morris	Date	18 MAR 07	File Name	GlobalWANBackbone.vsd

Global Routing

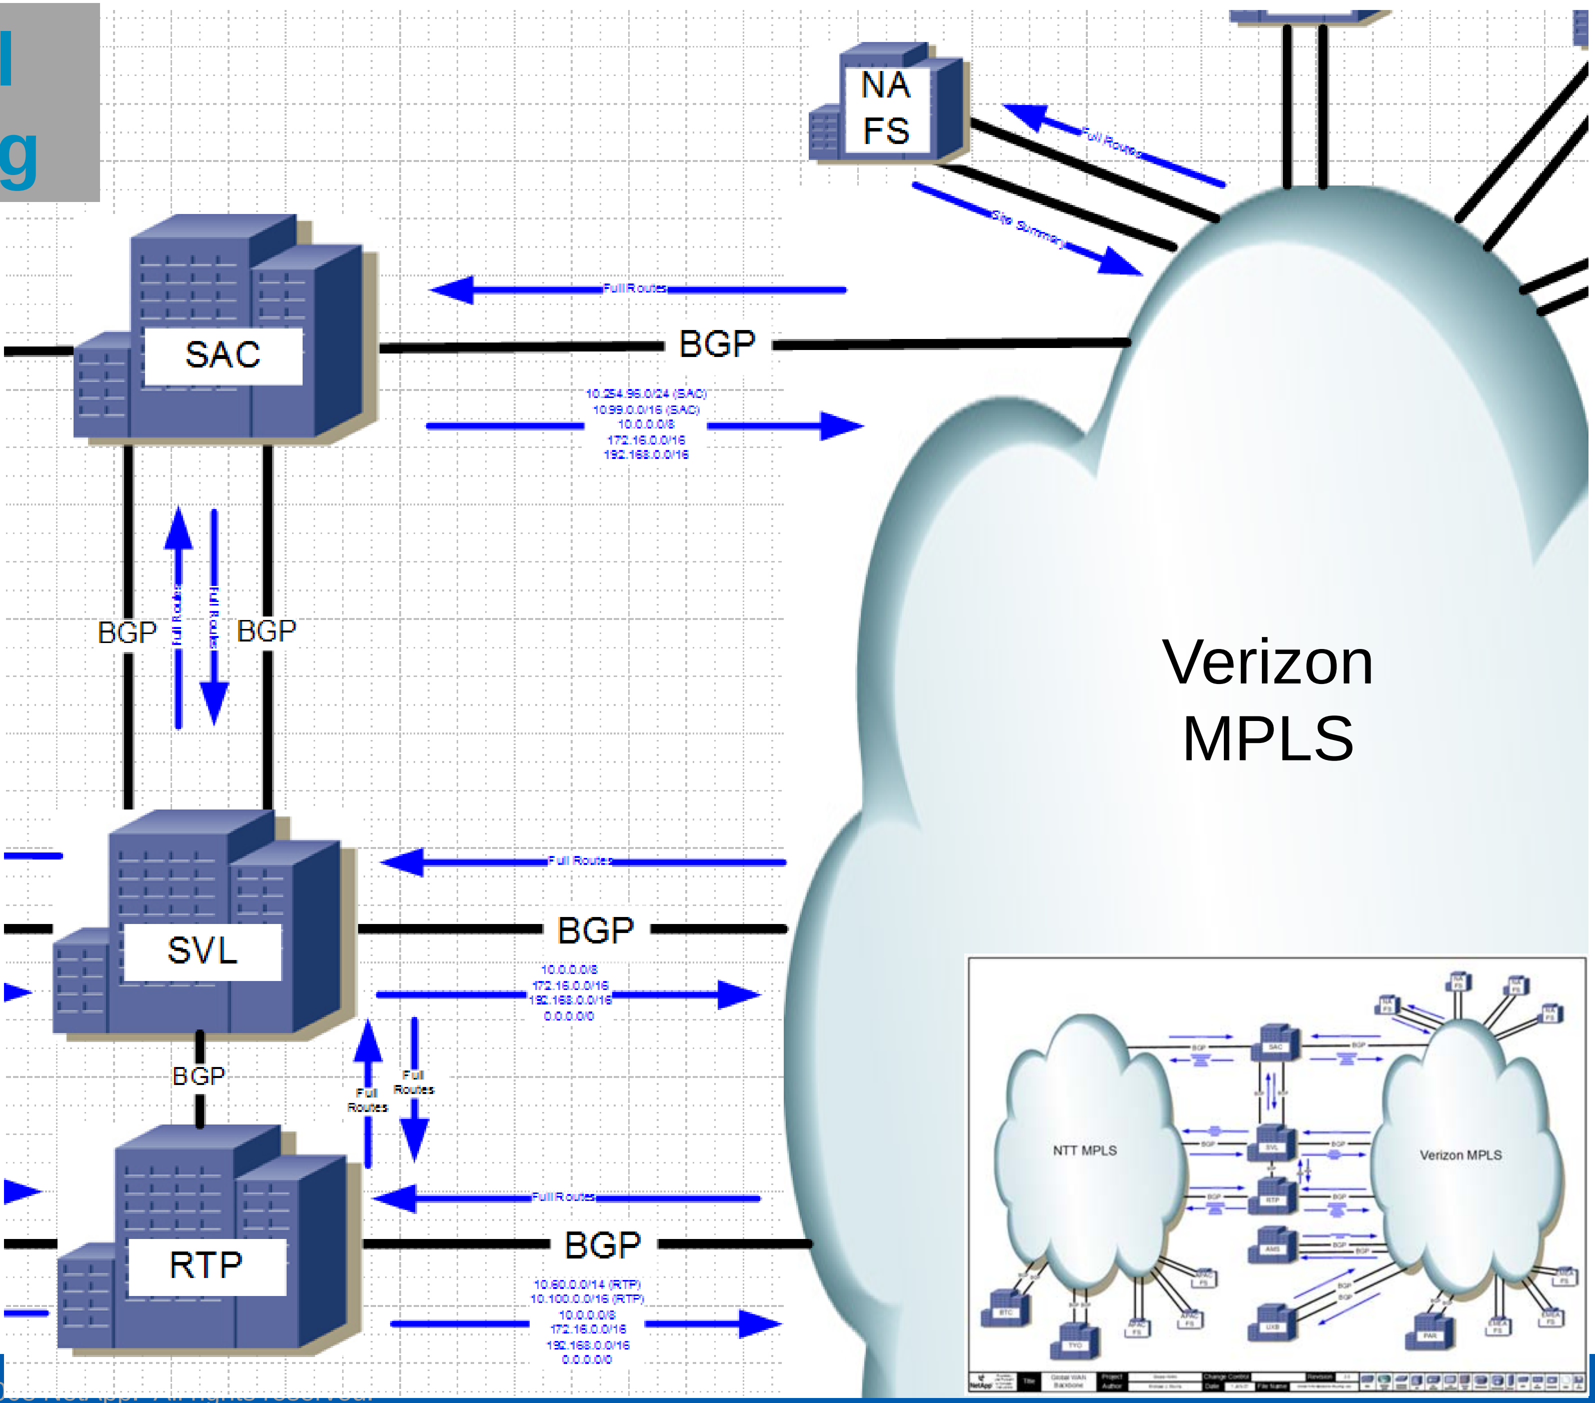

QoS

- 7-level QoS model
- Maps to any MPLS carrier QoS model by combining traffic classes.
- Protects business application and VoIP from bulk transfers and Internet traffic

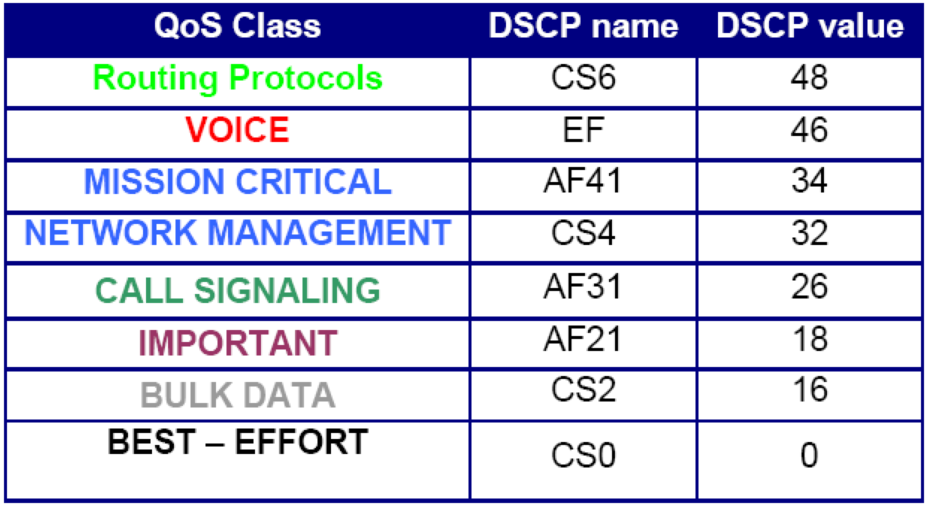

QoS Class	DSCP name	DSCP value
Routing Protocols	CS6	48
VOICE	EF	46
MISSION CRITICAL	AF41	34
NETWORK MANAGEMENT	CS4	32
CALL SIGNALING	AF31	26
IMPORTANT	AF21	18
BULK DATA	CS2	16
BEST – EFFORT	CS0	0