

## Designing Corporate Networks for VPLS Carriage

*Nextgen's National VPLS network offers a corporate "wide area network" that behaves like a LAN segment. In particular, it provides virtual, MAC-learning, Ethernet switches for each customer with advanced quality of service features, transparent carriage of corporate VLANs and high-speed interfaces.*

*This feature set has many advantages for corporate networks over the alternatives of ATM, Frame Relay and IP-VPN. Making most effective use of these features requires a little design consideration and possibly some network reconfiguration.*

*This paper offers a "best practises" design guide addressing the common issues customer's should consider when implementing their corporate network using Nextgen's National VPLS service.*

### Contents

1	Network Segmentation .....	2
2	Quality of Service.....	11
3	High Speed Networks .....	13
4	Security.....	16
5	Design for Migration.....	17
6	Conclusion .....	19
7	References .....	19
8	Acronyms.....	20
9	Contact .....	20

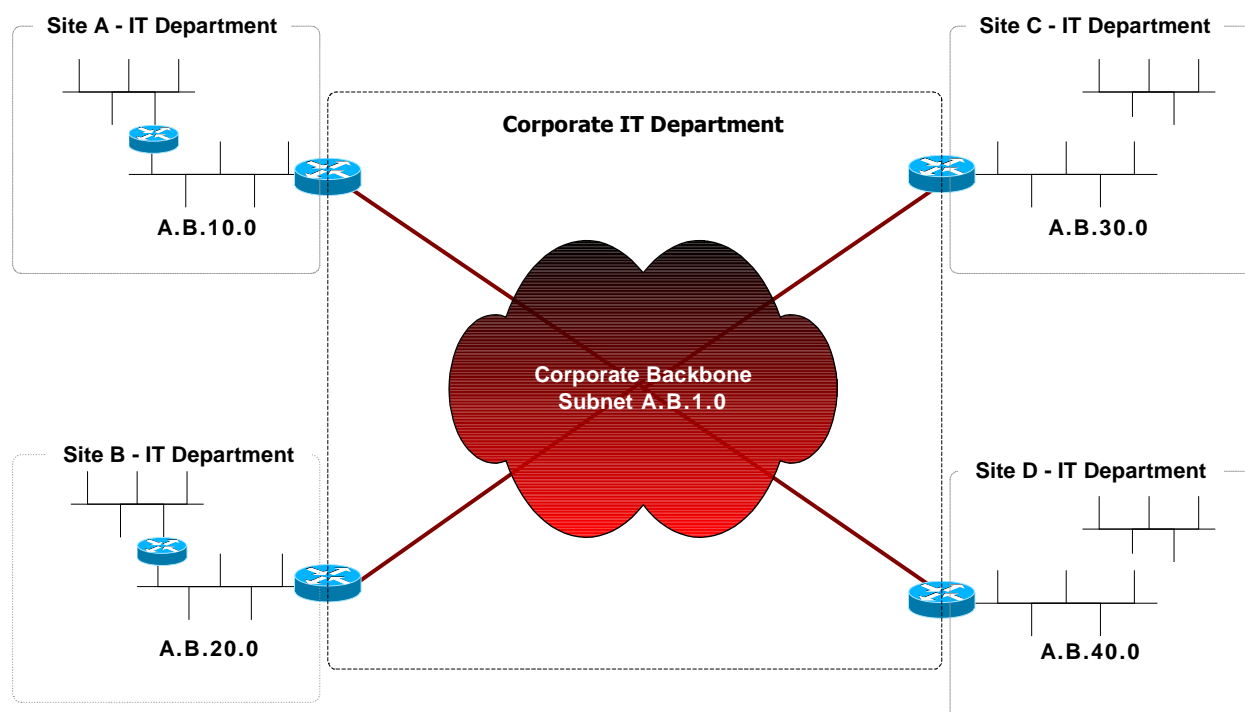
# 1 Network Segmentation

Nextgen's National VPLS network could be operated as a single, flat Ethernet segment, with intelligent Ethernet switches connecting every desktop to each other and to the Internet Gateway router. This design might be quite reasonable for a smaller organisation up to a few hundred IP clients.

Current "Best Practise" in corporate network design strongly recommends segmenting the corporate network into smaller, more manageable network chunks. These sub-networks are usually arranged according to who is responsible for administering the sub-network.

Breaking networks down into sub-networks based on administrative domain ensures that most network faults (failure of hardware or configuration) can usually resolved by a single responsible IT manager.

Historically, these sub-networks have involved dedicated network infrastructure and are attached to a corporate "backbone" by a router at each geographic location. The "backbone routers" are usually the responsibility of the corporate IT department. The local IT administrators in this model deal with desk-top, server and LAN communication issues, while any routing issues are escalated to the corporate IT department. This is shown in Figure 1.



**Figure 1: Administrative Responsibility**

This clean segmentation model breaks down where a single administrative domain covers multiple geographic locations. In this environment, a single IT administrator may be responsible for administering IT for a single department or business unit at two different sites, but where multiple departments or business units share the site. Examples might be the finance and engineering divisions, or retail and wholesale business units. Such an arrangement is more common for larger organisations and is shown in Figure 2.

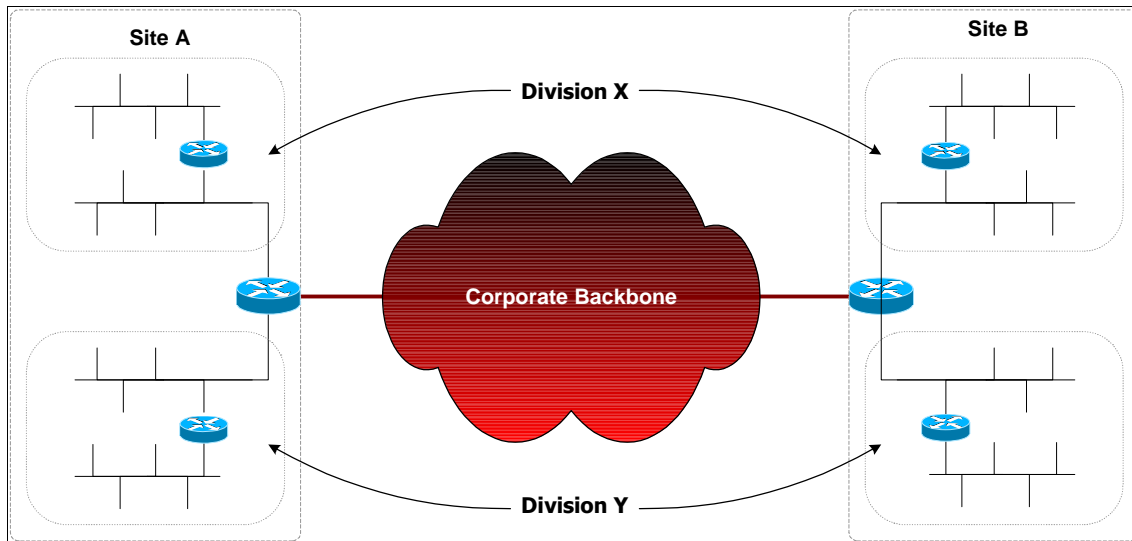


Figure 2: Administration across multiple geographic locations

## 1.1 Virtual Local Area Networks

A technology solution that helps the corporate management of disparate IT departments is to “trunk” segmented networks over a shared Ethernet switched backbone. The technology typically employed for this is defined in the technology standard IEEE 802.1Q [1] where an extra identifying tag is inserted near the front of each Ethernet frame: this technology is known as Virtual Local Area Network trunking “VLAN trunking”.

With VLANs, geographically diverse networks can be treated as a single logical (or virtual) local area network. The backbone can be based on Ethernet switches and the corporate routing network can be significantly simplified. This is shown in Figure 3A. Nextgen’s National VPLS network carries VLANs transparently, and so fully supports this type of network segmentation.

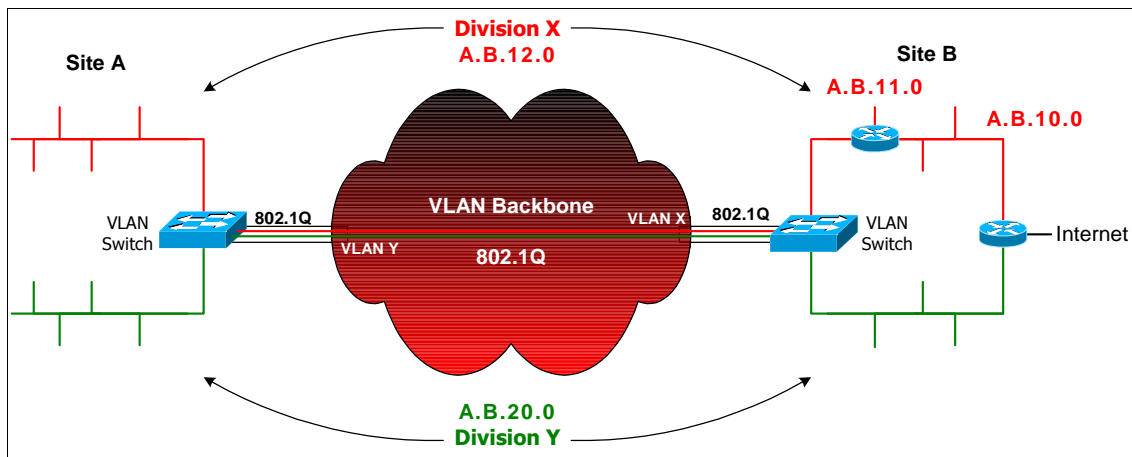


Figure 3A: Split Geography using VLANs for Administrative Domain Management

Typically, each virtual LAN is configured with its own IP sub-network, with a common physical interface being configured with multiple logical interfaces – one logical interface with a unique IP subnet identifier per VLAN. Traffic can flow between VLANs, but only through the IP router.

When configured like this, logically the network can be treated and administrated as though the virtual LAN segments are isolated. This is shown in Figure 3B.

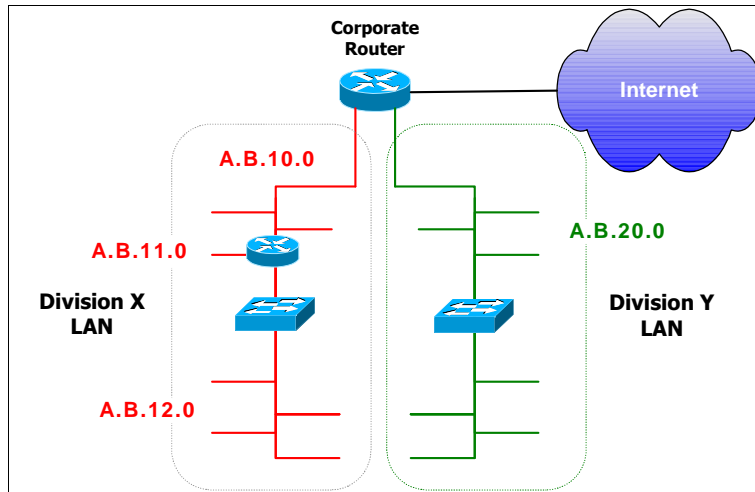


Figure 3B: Logical Network using VLANs

## 1.2 Virtualized WAN

In some cases, for instance large, conglomerate organisations, Virtual LAN concepts do not sufficiently support the corporate network’s natural administrative domains. For these cases, Nextgen offers a Virtualized WAN capability.

Virtualized WAN allows for an extra level of network segmentation, where multiple VPLS instances can be configured. Each VPLS instance supports multiple transparent VLANs. This is shown in Figure 4, where, for instance, each business unit or subsidiary company can operate on its own VPLS Segment with its engineering, finance, human relations, retail and wholesale specific VLANs. Typically a separate WAN instance is configured for the corporate IT department to provide high level network aggregation – for example, to allow bulk buying of services such as internet, web presence and email.




Figure 4 Virtual WAN solution

As with VLANs, each WAN instance is usually configured with its own IP subnet that is then further segmented into VLAN specific subnets. Traffic can flow between the various VLANs and WANs only by traversing well defined edge routers so that a set of hierarchical policies can be easily and independently maintained.

## 1.3 Segment Size Limits

There are a number of factors that can effect the optimal network segment size, including:

- Ethernet efficiency
- Routing convergence times
- Multicast behaviour

### 1.3.1 Ethernet Efficiency (1000 MAC Addresses)

Ethernet's great advantage is that it is an easy to configure technology that auto-discovers newly attached network devices by broadcasting address resolution packets.

For very large Ethernet networks (with a large number of IP-Clients), this broadcasting behaviour can cause the WAN network efficiency to degrade, even if the smartest of MAC learning switches are deployed. Optimal performance is difficult to determine, and networks can be tuned for better performance using MAC learning parameters.

A "rule of thumb" can be determined based on the typical MAC time-out of 5 minutes (IEEE 802.1D default MAC Aging of 300.0 seconds [2]). The network load introduced by having to flood a maximum (1522Byte) packet every time a MAC address is aged out is roughly 40bits/second for each IP client. At 1000 clients, this could consume as much as 10% of a lowest rate (512kbit/s) VPLS access circuit.

While Nextgen's VPLS network can handle much larger networks, for efficiency reasons, we recommend no more than 1000 MAC addresses be visible to each VPLS instance.

### 1.3.2 Routing Convergence Time (100 IP Routers)

A commonly quoted limit to the size of a network is given by the performance of interior routing protocols, typically router memory and CPU consumption.

OSPF, a common interior gateway protocol, is shown in IETF-RFC2329<sup>1</sup> [3] to be deployed with typically 100 routers in each OSPF area (up to 350) and typically 23 areas (up to 60) in a domain. An OSPF Area is a local organisation of routers that communicate directly with each other but elect to communicate with the rest of the world via nominated representatives (Designated Routers).


Cisco's OSPF design guide [4] gives a number of considerations for maximum area sizes:

- What kind of OSPF Area do you have?
- What kind of CPU power do you have?
- How many external LSAs?
- Are you running NBMA networks?
- How well are areas summarized?

---

<sup>1</sup> Internet Engineering Task Force – Request for Comment number 2329. Request for Comment are the Internet Standards documents, but include many other documents including observations, practise guides and even jokes.

While clearly very large OSPF networks can be deployed, a common design guide is to design for fewer than 100 routers per OSPF area. Good practise based on these sorts of numbers is to limit the number of meshed routers in an area to be around 50: a good compromise for manageability while allowing some room for unplanned growth. These areas can each be located on a separate VLAN within one or more VPLS clouds as shown in Figure 5.



**Figure 5: Area Design with VLAN and VWAN**

The designated router (DR) and backup designated router (BDR) are logical locations for Area Border Routers that connect the specific areas onto the backbone area (Area 0) as they are more directly under the control of the central IT group and likely to be more powerful routers. These border routers link between VLANs on a single interface and can link between VPLS instances on physical interfaces where virtualised WAN's are used.

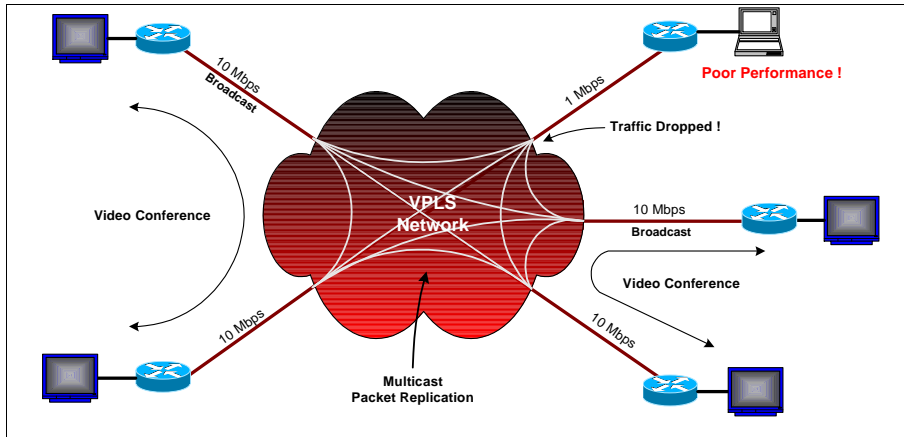
To further optimise the routing network, stubby areas<sup>2</sup> should be defined and default routes used. These techniques can increase both the effective area size and effective network size supported on a single VPLS instance.

## 1.4 Multicast and Broadcast

Apart from the address resolution broadcasts and the flooded packets where destinations are not known, some streaming traffic is configured to broadcast throughout the network. Where this broadcast traffic is “trimmed” to those devices requesting the traffic, the traffic is known as multicast. A very common method for trimming multicast traffic is to perform IGMP Snooping [5].

As with standard, default Ethernet behaviour, any broadcast and multicast packets sent into the VPLS cloud are replicated and sent out of every interface. This can result in service degradation for low rate access circuits that are not interested in the multicast traffic. Multicast Ethernet networks should be designed to ensure that low-rate access circuits are not unduly affected by high rate multicast traffic. Figure 6 indicates the risk with low rate interfaces.

<sup>2</sup> Stubby Areas are defined in the OSPF to represent sub-networks where all external traffic is directed through a small number of routers and most routers can be hard configured with a default next hop.



**Figure 6: Broadcast Flooding**

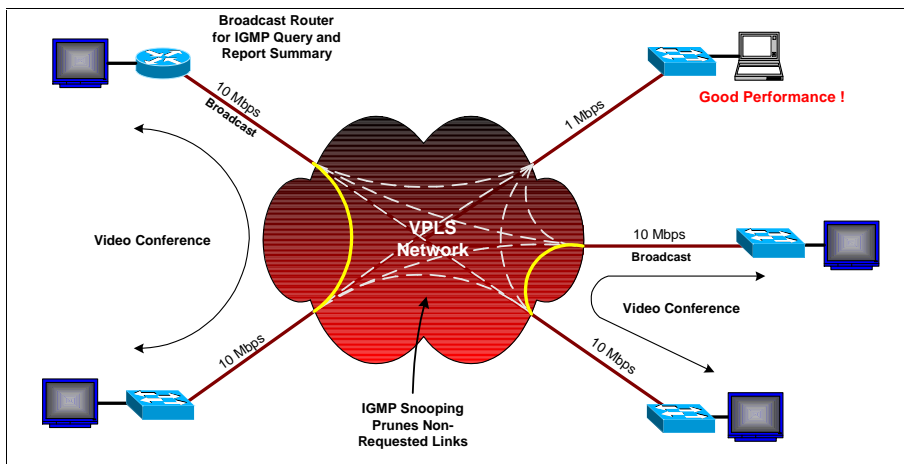
Since the Nextgen VPLS instance is transparent to VLAN<sup>3</sup>, it is not sufficient to segment this multicast traffic using VLANs. Nextgen offers two alternatives to this.

- IGMP Snooping
- Virtual WAN

**1.4.1 IGMP Snooping Solution**

IGMP Snooping is where an Ethernet switch monitors the IGMP queries and reports and keeps a multicast-forwarding table so that only interfaces with a demonstrated interest in a multicast group receive that traffic. With IGMP snooping, no signalling traffic is modified so that clients and querying servers are unaware of the efficient forwarding of the traffic.

Nextgen’s National VPLS network can perform IGMP snooping, which can act to alleviate the broadcast problem. The benefits of IGMP snooping are shown in Figure 7.



**Figure 7: IGMP Snooping**

<sup>3</sup> That is, the VPLS network ignores the VLAN tag and does not participate in automatic VLAN configuration.

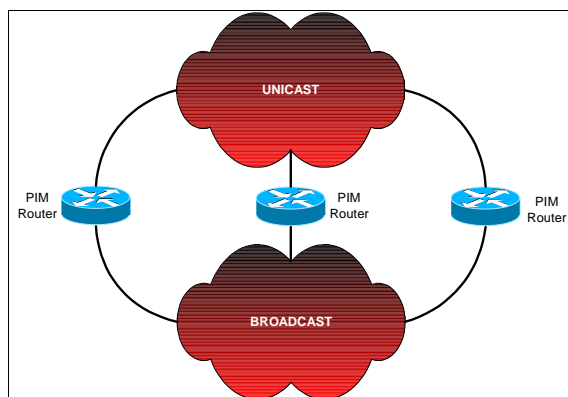
IGMP snooping only works when IGMP is used to signal the multicast traffic. Typically, IGMP is used between a small number of source routers and IP clients that are located on the same logical LAN segment. This means that IGMP snooping is of most benefit in the VPLS instance where a small number of routers are attached and the majority of access circuits are Ethernet switch based.

For signalling multicast between routers the usual protocol of choice is Protocol Independent Multicast (PIM) [6] and not IGMP. Nextgen are evaluating the option of performing the analogous PIM-snooping, but as yet have no firm plans for implementation of this feature.

#### 1.4.2 Virtual WAN Solution

If an alternative method of supporting multicast is chosen, such as PIM, then Nextgen can provide an isolated VPLS instance to support this kind of traffic. The “unicast” VPLS instance is used to connect all routers, while the multicast VPLS instance is only connected to those sites that have been dimensioned appropriately to handle the corporate multicast traffic. Virtual WAN is required since the VPLS network carries the VLAN tags without inspection.

A network provisioned with dedicated unicast and dedicated multicast VPLS instances could appear as shown in Figure 8.



**Figure 8: Broadcast V-WAN**

### 1.5 Best Practise Network Segmentation

- Segment Networks by Administrative IT Domain.
- Secondary segmentation may be by project, by functional group or by traffic flow.
- Use VLAN's where Administrative IT Domains are geographically spread
- Use Virtual WAN where more structured, hierarchical segmentation is required
- Each Segment should have a unique IP subnet
- Avoid putting multiple IP subnets on a single logical LAN
- Limit the size of Ethernet switched networks to less than a thousand IP clients
- Limit the size of interior routed areas to less than 100 routers
- Use area border routers to connect network segments, whether VLAN's or VPLS instances
- Use stubby routing networks where practical

- Consider placing multicast on an isolated VPLS instance
- Consider a separate high-speed national VPLS instance as a corporate backbone.

Two examples of possible network architectures are provided in Figure 9 and Figure 10.

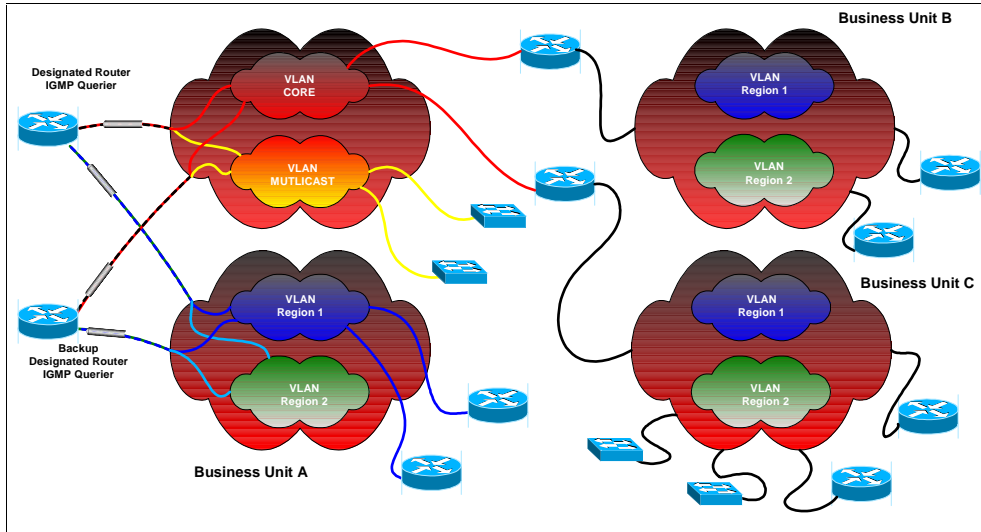


Figure 9: Example Network Design

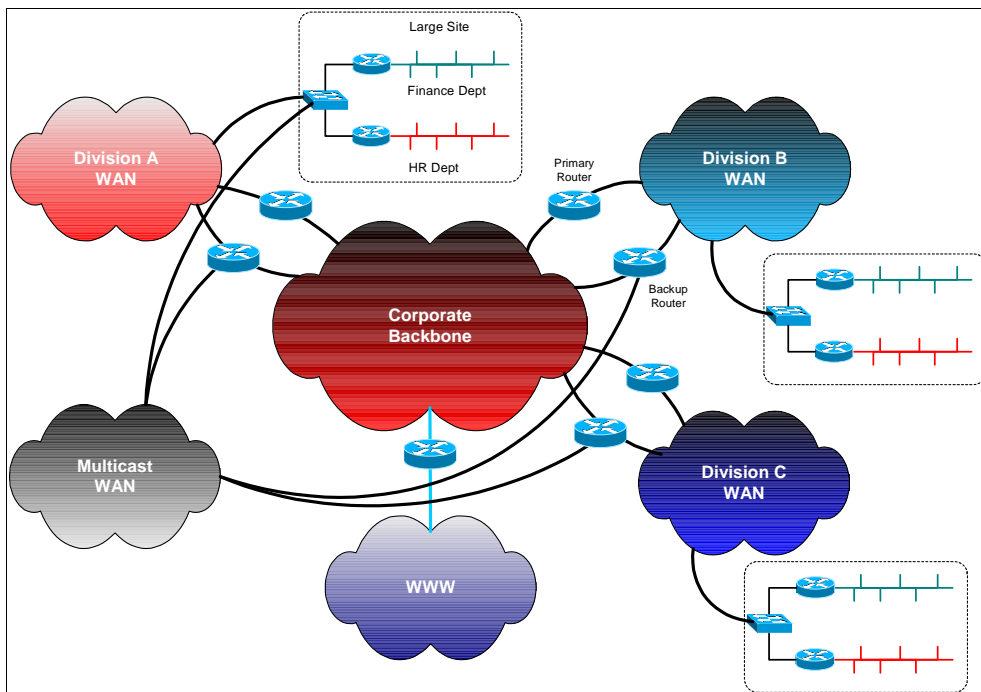


Figure 10: Example Network Design

## 2 Quality of Service

There are two schools of thought on how to correctly design networks for quality of service.

- Excessive Bandwidth
- Priority Queuing

### 2.1 Excessive Bandwidth

The philosophy of Excessive Bandwidth is that when bandwidth always exceeds demand, queues are always near empty and all packets receive lowest latency and jitter.

This works extremely well for higher speed interfaces, such as those used on the Nextgen VPLS core. As access-circuit bandwidths drop and the number of access circuits increases, the Excessive Bandwidth philosophy becomes compromised.

The maximum jitter experienced by a network with Excessive Bandwidth at an egress point is the number of access-circuits times the maximum packet size divided by the egress access circuit rate. A maximum jitter of around 30 milliseconds noticeably impacts voice services; five access-circuits at 2Mbit/s can generate 30 milliseconds jitter.

The basis for the “Excessive Bandwidth” solution is given in Figure 11, where the expected queue length for random arrivals and departures is given as a function of load, and the queue response to simultaneous packet arrivals indicated.

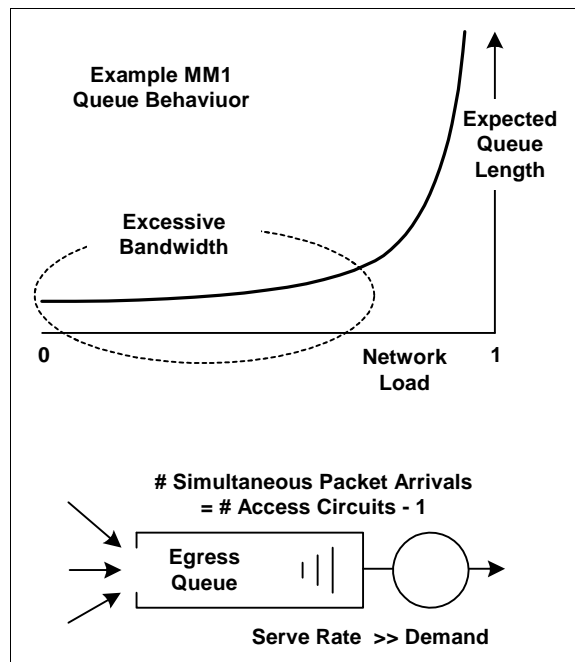


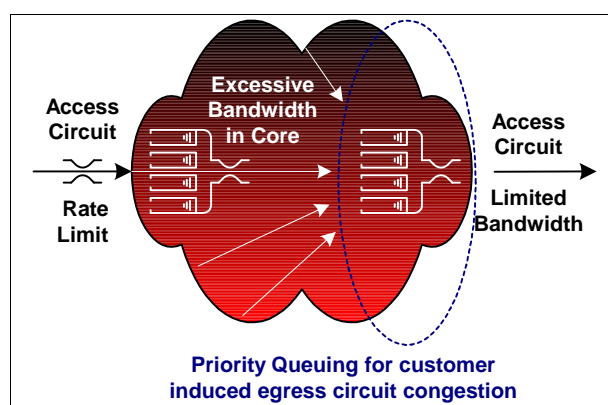
Figure 11: QOS - Excessive Bandwidth and Priority Queuing

## 2.2 Priority Queuing

Where excessive bandwidth fails, networks need to support differentiated quality of service, typically priority queuing. This is most likely to be needed where multiple access circuits concurrently transmit data towards a single lower rate access circuit.

Nextgen's National VPLS network provides hierarchical ingress buffering and shaping and hierarchical egress buffering and shaping. Four queues are provided at every interface for customer traffic and traffic is served from these queues in strict priority order. This is shown in Figure 12.

While fully supporting differentiated services, Nextgen's VPLS core network effectively operates on the "Excessive Bandwidth" paradigm. This helps to isolate customers from each other. In the core, priority queuing ensures sensible behaviour during a network failure that causes core network overload.



**Figure 12: QOS - Priority Queuing**

Nextgen's VPLS is configured to honour any classification applied by the customer. To make use of this, customers must configure their corporate network to correctly label all packets flowing into the VPN, using 802.1P [7], DSCP [8] or TOS [9] as agreed.

Good practise is for corporate networks to use the highest quality of service for network control traffic, the second quality of service for low latency traffic (typically voice and/or video), the third quality for high priority corporate data and the fourth for normal priority data.

Nextgen's VPLS network places traffic rate limits on all classes of service to help ensure desirable network behaviour. Best practise is for corporate network equipment to shape each traffic class to meet the policing criteria of the WAN network.

While Nextgen provides buffering and shaping at ingress to the VPLS core, this reduces customer visibility and should be treated as a safety for mis-configured corporate equipment.

Corporate networks should monitor their own network usage to establish baselines, facilitate trouble-shooting and plan for network upgrade. Best practise is to monitor peak and average five-minute loads at each class of service. A particularly sensitive indicator of increasing traffic is the maximum egress queue-length of each class of service. The monitoring points and a typical time development of traffic is shown in Figure 13.

Queue theory shows that as average demand increases, queue lengths increase hyperbolically. Good practise keeps traffic demand below 80% for good human response times and below 50% for low latency and jitter.

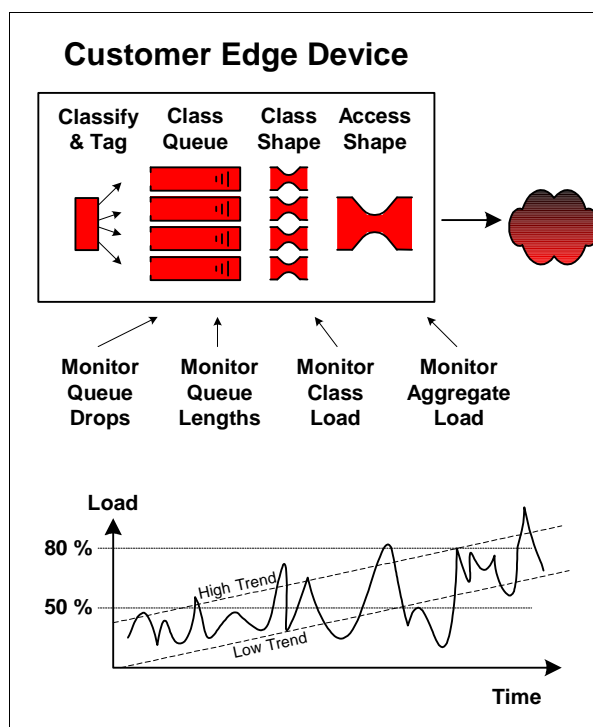


Figure 13: QOS - Best Practice

## 2.3 Best Practise QOS

- Match quality to business needs
- Tag all traffic into the WAN
- Shape traffic to match WAN policing
- Monitor both ingress and egress traffic loads
- Monitor peak and average 5 minute loads
- Monitor loads for each traffic class
- Monitor peak queue lengths for each traffic class
- For responsive traffic keep average load below 80%
- For voice and video traffic keep average load below 50%

## 3 High Speed Networks

High-speed networks with national coverage introduce some protocol design considerations. These networks have a high bandwidth-delay product and are known in the literature as “long fat pipes” [10].

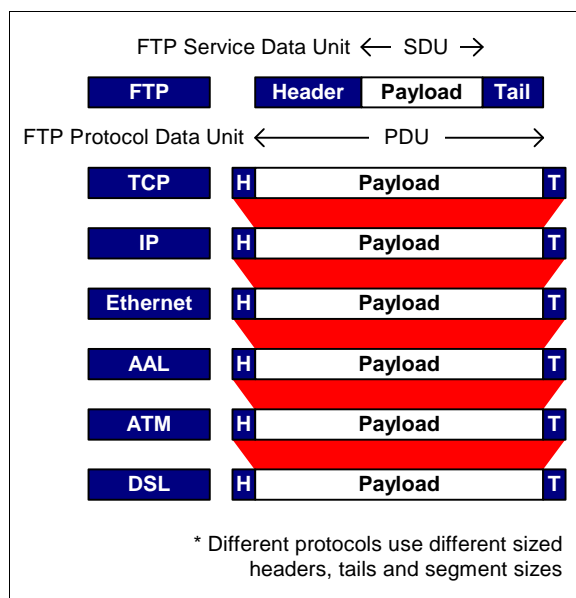
These “long fat pipes” often result in common throughput performance tests indicating lower network throughput than actually available. To avoid these traps, RFC 2544 [11] testing is recommended.

A typical method of determining network performance is an FTP download or upload. Unfortunately, TCP<sup>4</sup> has a number of tuning parameters that are by default not suitable for high bandwidth-delay product networks.

<sup>4</sup> TCP is the underlying transport layer protocol for FTP.

## 3.1 Incorrect Protocol Data Unit

FTP typically reports the FTP payload data rate (PDU), while WAN networks typically provide either a layer-2 service data rate (SDU) or a layer-2 PDU rate. There can be significant difference between these rates due to the various overheads encountered in encapsulation. The various protocol layers are shown in Figure 14, as a guide. The specific overheads tend to vary by protocol and protocol options selected.




**Figure 14: Protocol Overheads**

Where clarity of data rate is required, the protocol layer and specific frame points being measured need to be described, so that an “Ethernet Rate” is not as specific as “Ethernet PDU including MAC and FCS, but ignoring preamble and start byte”.

## 3.2 TCP Slow Start

TCP has an exponential ramp-up, known as slow-start [12]. For a high bandwidth delay product, a significant portion of the FTP file is transmitted during this start-up period and the measured throughput will be lower than the available network throughput. The dependence on file size for an accurately reported FTP data rate test is shown in Figure 15.



**Figure 15: TCP - Slow Start**

If FTP throughput is used, for example where a quick convenient test where error is not important, increasing the file size can provide a better estimate of the throughput. Better is to use a method designed to measure peak throughput without relying on TCP hunting behaviour.

### 3.3 Max Window Size

The TCP maximum window size limits the unacknowledged “bits in flight” [10]. For optimal performance, the TCP maximum window size should be greater than the network bandwidth times the round-trip-delay [13].

As an example, 10Mbit/s and 50ms RTT gives an optimal TCP Max Window Size of greater than 50kBytes. For comparison:

- Windows 98 = 2kBytes
- Windows 2000 = 17.5kBytes
- Windows XP = 64kBytes

The maximum TCP throughput is limited to the maximum number of bits in flight every round trip time. For a 50ms RTT, maximum throughput is:

- Windows 98 = 320kbit/s
- Windows 2000 = 2.8Mbit/s
- Windows XP = 10.2Mbit/s

An example of the impact TCP Maximum Window Size may have on the transmit patterns, and hence achieved throughput is shown in Figure 16.

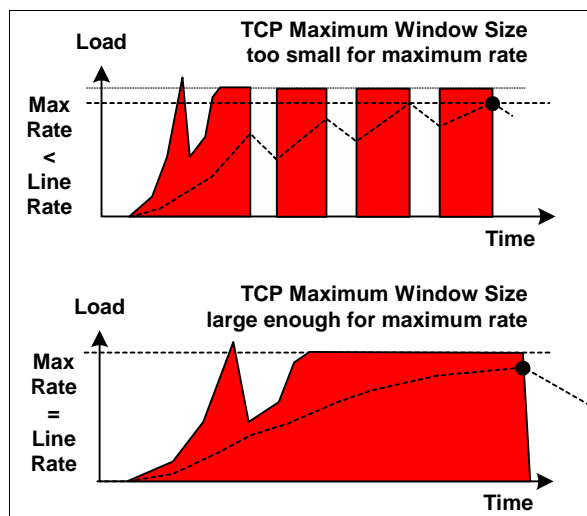


Figure 16: TCP - Maximum Window Size

### 3.4 Best Practise Performance Validation

When validating the performance of a network, best practise is:

- Use RFC 2544 to measure network throughput
- Know your protocol overheads
- Specify the measurement points clearly.

If FTP must be used, then

- Use large files to limit error
- Optimise TCP configuration to limit error
- Try to use this style of test as a quick indicative measure only
- Allow for FTP under-reporting the actual throughput

## 4 Security

Security of traffic can mean many things:

- Inability for a third party to intercept data
- Inability of a third party to modify data
- Inability for a third party to consume network resources (theft)
- Inability for a third party to consume network resources (denial of service)

Nextgen's VPLS network provides a secure private VPN that is isolated from the Internet. Being based on MPLS (a connection oriented network protocol) it is as secure as any other VPN offering.

The reality is that carriers are always able to intercept traffic and attack any encrypted traffic at their leisure – in fact, the ability to do this is a legal requirement in most countries [14]. Using current encryption techniques, it is purely a matter of time before the information is exposed.

If true network security is required, the customer should be using secured, point-to-point tunnels with three way handshaking and secrets contained entirely within the trusted organisation. The encryption key should be as long as possible (256bit or better) and should be changed regularly.

Any security equipment should be sourced from reputable and trusted equipment providers.

## 4.1 Best Practise Security

- Seek a private VPN from a reputable supplier.
- Use encrypted tunnels with strong and frequently changing encryption keys.
- Manage your own shared-secrets for key generation.

## 5 Design for Migration

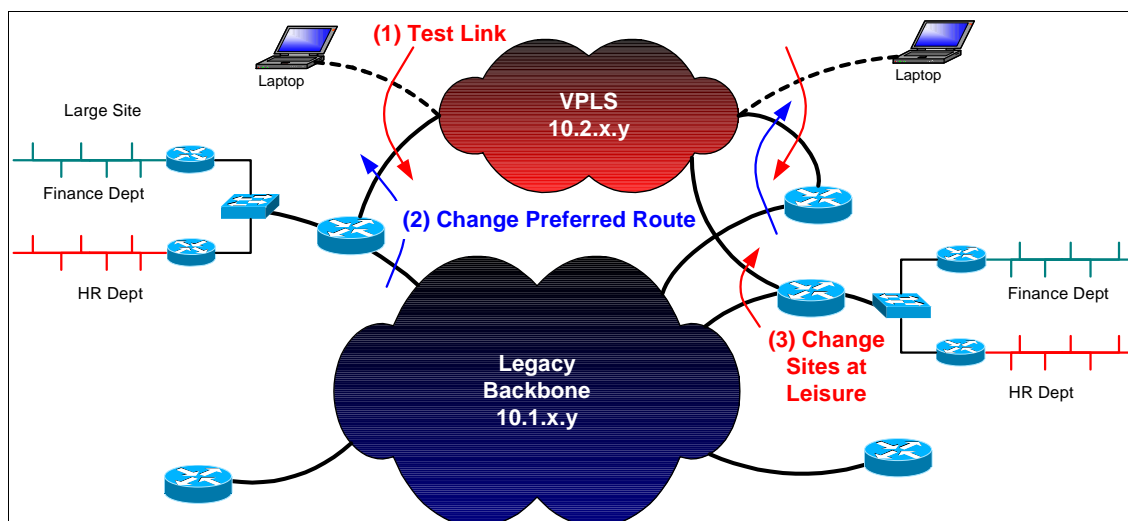
Successful network migration is most easily achieved with careful consideration and planning. Factors to consider include:

- Current and future needs of the corporate network
- Resources required to migrate
- Acceptance procedures, including acceptance criteria, pilot networks and proving-in periods
- Training requirements
- Rollback procedures in case of failed migration

Nextgen's National VPLS is a product that minimises the complexity of migration and is relatively easy and convenient to install. Being a layer-2 VPN, it will appear as a single logical LAN segment presenting at an RJ45 socket – or multiple logical LAN segments if VLAN's and/or multiple VPLS Instances are employed. Any equipment connected to this socket will have Ethernet visibility of any other equipment connected to the VPLS instance.

It is possible to seamlessly cutover from a legacy VPN to Nextgen's VPLS during business hours. The following process is effective and also shown in Figure 17:

- Test service connectivity – requires two services to be connected.
- Allocate an IP subnet (or multiple if using multiple VLANs) for the VPLS cloud exactly as any other Ethernet subnet would be configured.
- Configure a network equipment interface – typically but not necessarily on a gateway router – for operation on this subnet with a priority weighting (usually lower priority) so that the interface operates as a backup to the legacy network. Connect the configured interface to the VPLS and allow the router to discover its neighbours.
- Verify that the various neighbours can see each other.
- Modify the routing priorities to make the VPLS the primary network. (One option to consider here is to administratively disable the legacy network interface – a quick and easily reversed process). Traffic should migrate on to the VPLS cloud.
- Disconnect the legacy network once satisfied that everything is working correctly.



**Figure 17: Network Migration**

It is good practise to migrate small segments of the network at a time to minimise the risk involved and to establish experience and confidence with the process. The segmented network design provides a good guide for the logical migration entities.

Nextgen recommends a central router as one of the first activation sites. This router remains connected to both VPLS and Legacy network as a gateway to bridge the two halves of the network for the duration of the migration program.

For large corporations employing multiple VPLS instances, the various instances can be migrated in isolation. A sensible process is to fully migrate a smaller VPLS instance and then migrate the backbone VPLS instance before proceeding to the remaining VPLS instance. This maximises both the utility of VPLS and the corporate IT departments experience with migration to VPLS before tackling the generally harder coordination between corporate and divisional IT departments.

## 5.1 Best Practise Migration

- Plan a staged migration starting with a smaller network segment
- Use the target network model as a guide for migration planning
- Ensure a roll-back plan and sufficient resources are in place before starting
- Include a more central router as an early connection to the VPLS network
- Leave the central router(s) connected to the legacy VPN until migration of all sites is completed
- Use the VPLS network and legacy VPN as redundant networks to minimise risk and provide for easy roll-back

## 6 Conclusion

Nextgen's National VPLS network offers a corporate "wide area network" with superior features to many legacy WAN technologies.

While designing the corporate WAN as a simple flat Ethernet segment is entirely possible and may be ideally suited for a smaller enterprise network, the rich features of the Nextgen VPLS with its high rates, transparent VLAN support, integrated QoS and multiple virtual WAN capabilities make it equally suited for the most advanced and organisationally complex corporation.

Making most use of the Nextgen VPLS feature set requires a little planning and design work but the resulting solution will be aligned with corporate best practice networking.

Migration from an existing wide area solution can be performed with minimal impact and staged to match available resources.

## 7 References

- [1] IEEE Std 802.1Q-2005 "IEEE Standard for Local and metropolitan area networks - Virtual Bridged Local Area Networks.", IEEE, New York, 2006
- [2] IEEE Std 802.1D-2004 "IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges, IEEE, New York, 2004
- [3] Moy, J, "IETF RFC 2329: OSPF Standardization Report", IETF, April 1998.
- [4] Cisco Systems, "OSPF Design Guide", Cisco White Paper, Document-id 7039, April 2006.
- [5] Christensen, M et. al., "IETF RFC 4541: Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches", IETF, May 2006
- [6] Fenner, B. et al, "IETF RFC 4601: Protocol Independent Multicast – Sparse Mode (PIM-SM): Protocol Specification (Revised)", IETF, August 2006
- [7] IEEE 802.1P "IEEE Standard for Local and Metropolitan Area Networks: LAN Layer 2 QoS/CoS Protocol for Traffic Prioritization." In IEEE 802.1Q and IEEE 802.1D
- [8] Black, D et al, "IETF RFC 3140: Per Hop Behavior Identification Codes", IETF, June 2001
- [9] Postel (ed) "IETF RFC 791: Internet Protocol Specification", IETF, September 1981
- [10] Jacobson, V., Braden, R., "IETF RFC 1072: TCP Extensions for Long-Delay Paths", IETF, October 1988
- [11] Bradner, Z, McQuaid, J, "IETF RFC 2544: Benchmarking Methodology for Network Interconnect Devices", IETF, March 1999
- [12] Allman, M., Paxson, V., Stevens, W., "IETF RFC 2581: TCP Congestion Control", IETF, April 1999
- [13] Floyd, S., "IETF RFC 3649: HighSpeed TCP for Large Congestion Windows", IETF, December 2003
- [14] Telecommunications (Interception and Access) Act 1979 – Act Number 114 of 1979, Compilation of 30 December 2006 incorporating amendments up to Act No 86 of 2006. Attorney-General's Department, Canberra.

## 8 Acronyms

ATM	Asynchronous Transfer Mode
BDR	Backup Designated Router
CPU	Central Processor Unit
DR	Designated Router
DSCP	Diff Serve Code Point
FCS	Frame Check Sum
FTP	File Transfer Protocol
IEEE	Institute of Electrical and Electronic Engineers
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IT	Information Technology
LAN	Local Area Network
MAC	Media Access Control
Mbit/s	1,000,000 bits per second
MPLS	Multi Protocol Label Switching
OSPF	Open Shortest Path First
PDU	Protocol Data Unit
PIM	Protocol Independent Multicast
QOS	Quality of Service
RJ-45	Registered Jack 45 – more correctly 8P8C
RTT	Round Trip Time
SDU	Service Data Unit
TCP	Transmission Control Protocol
TOS	Type Of Service
VLAN	Virtual Local Area Network
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
WAN	Wide Area Network

## 9 Contact

Written and published by Nextgen Networks Pty Limited, January 2007.

**Nextgen Networks Pty Limited**  
**PO Box 13071**  
**Law Courts Post Office**  
**Melbourne Vic 8010**

**1300 653 351**