

PSN - Technical Domain Description



**Making
government
work better**

Document Details

Document Name	PSN – Technical Domain Description
Creation Date	15 th March 2010
Author	Nick Higgins
Version	1.0
Status	Final

Version History

Version 1.0	TDD version completed with Industry and Government comment.
-------------	---

Document Approval

Name	Position	DATE
Geoff Hall	Design Authority Chair	15 th March 2010
Martin Gould	Programme Manager	15 th March 2010
John Stubley	Programme Director	15 th March 2010

Table of Contents

1	DOCUMENT SCOPE	5
1.1.1	<i>Document Relationship to the PSN Operating Model</i>	5
2	MANAGEMENT OF CONSTRAINED LOGICAL RESOURCES IN THE PSN	6
2.1	PUBLIC LOGICAL RESOURCES.....	6
2.2	PRIVATE LOGICAL RESOURCES.....	6
2.3	INSTANCES OF PRIVATE LOGICAL RESOURCES WITHIN THE PSN STANDARDS.....	6
2.4	DNS SERVICES AND INTER-IMPACT LEVEL (IL) DOMAIN BROWSING	6
2.4.1	<i>Current Public Sector Scenarios</i>	6
2.4.2	<i>PSN Hierarchical DNS Model</i>	7
2.4.3	<i>Inter Impact Level Domain Recursion</i>	7
2.4.4	<i>Authoritative DNS Root Servers and the PSN</i>	7
2.4.5	<i>Domain Naming Conventions and Explicit Impact Level Identifiers</i>	7
2.4.6	<i>PSN Service Instances Impacted by DNS and Domain Management</i>	7
3	GCN – INTER PROVIDER MPLS	8
3.1	ASSUMPTIONS	8
3.2	REFERENCE DOCUMENTS	8
3.2.1	<i>RFC Notes</i>	9
4	THE GOVERNMENT CONVEYANCE NETWORK (GCN)	9
4.1	GCN - CONCEPTUAL ARCHITECTURE	9
5	GCN INTER-PROVIDER MPLS	10
5.1	INTER AUTONOMOUS SYSTEM (INTER-AS) MPLS	11
5.2	OPTION A: BACK- TO-BACK VRF.....	11
5.2.1	<i>Packet Forwarding</i>	11
5.2.2	<i>Security</i>	14
5.2.3	<i>Scaling</i>	14
5.3	OPTION B: ASBR-TO-ASBR MULTI-PROTOCOL eBGP FOR VPNV4.....	14
5.3.1	<i>Packet Forwarding</i>	15
5.3.2	<i>Security</i>	17
5.3.3	<i>Scaling</i>	17
5.4	QUALITY OF SERVICE (QOS)	17
6	GCN INTER-AS DEPLOYMENT MODEL	18
6.1	BASELINE ARCHITECTURE	18
6.1.1.1	Scenario 1: A/A/A	20
6.1.1.2	Scenario 2: A/B/A	20
6.1.1.3	Scenario 3: B/B/B	21
6.1.1.4	Scenario 4: B/A/B	22
6.1.1.5	Scenario 5: Both A/B Interconnect Points	22
7	GCN NETWORK TO NETWORK INTERFACE (NNI) - COMMON ATTRIBUTES	23
7.1	PHYSICAL DEPLOYMENT.....	23
7.1.1	<i>Device Location and Security</i>	23
7.1.2	<i>Interface Allocation</i>	24
7.1.3	<i>Cabling</i>	24
7.1.4	<i>Labelling</i>	24
7.1.5	<i>Link Failure</i>	24
7.1.6	<i>Demarcation</i>	24
7.2	LAYER 2.....	24
7.2.1	<i>Virtual LAN (VLAN) Allocation</i>	24
7.3	LAYER 3.....	24
7.3.1	<i>IP Addressing</i>	24
7.4	ROUTING PROTOCOLS	24

7.4.1	Peering	25
7.4.2	AS Numbering and AS Path	25
7.4.2.1	Support for 32 bit AS Number Plan	25
7.4.3	Timers.....	25
7.4.4	Authentication.....	25
7.4.5	Prefix Limits.....	25
7.4.6	Attribute Support	25
7.4.7	Filtering	25
7.4.8	Load Balancing.....	25
7.4.9	VPN Configuration - Option A	25
7.4.10	VPN Configuration - Option B	26
7.4.10.1	RD Value	26
7.4.10.2	RT Re-write	26
7.4.11	VPN Option A and B Support.....	26
7.4.12	Scaling.....	26
7.4.13	Attack Mitigation.....	26
7.4.14	Monitoring.....	26
8	PSN TELEPHONY STANDARDS	26
8.1.1	PSN Telephony Standards - Reference Documents	27
8.2	APPLICATION OF STANDARDS.....	27
8.3	USE OF A MEDIATION SERVICE	28
8.4	INTERCONNECT MODEL	29
8.5	BASIC CALL	29
8.6	INTEROPERATION STANDARDS.....	29
8.6.1	Location Identifier	29
8.6.2	CODEC Selection.....	29
8.6.3	Negotiation	30
8.6.4	Transport Protocol	30
8.6.5	Network Authentication.....	30
8.6.6	Real Time Protocol (RTP) Transport Media	30
8.6.7	Voice Band Services.....	30
8.6.8	Calling Line Identity (CLI).....	30
8.6.9	Numbering Plan	30
9	APPENDIX A – TECHNICAL ABBREVIATIONS	31

Table of Figures

FIGURE 1: GCN – CONCEPTUAL ARCHITECTURE.....	9
FIGURE 2: BACK-TO-BACK VRF – CONNECTIVITY	11
FIGURE 3: BACK-TO-BACK VRF – CONTROL AND DATA PLANE	13
FIGURE 4: ASBR-TO-ASBR – CONNECTIVITY	15
FIGURE 5: ASBR-TO-ASBR – CONTROL AND DATA PLANE	16
FIGURE 6: GCN – GENERIC BASELINE ARCHITECTURE	19
FIGURE 7: SCENARIO 1 – MPLS OPTION A/A/A	20
FIGURE 8: SCENARIO 2 – MPLS OPTION A/B/A	21
FIGURE 9: SCENARIO 3 – MPLS OPTION B/B/B.....	21
FIGURE 10: SCENARIO 4 – MPLS OPTION B/A/B	22
FIGURE 11: SCENARIO 5 – BOTH MPLS OPTION A/B INTERCONNECTS	23
FIGURE 12: PSN TELEPHONY – MULTIPLE DEPARTMENTS USING MULTIPLE SERVICES	27
FIGURE 13: PSN TELEPHONY – PROVIDER INTERCONNECT.....	28
FIGURE 14: PSN TELEPHONY – MEDIATION SERVICE	28
FIGURE 15: PSN TELEPHONY – LOGICAL INTERCONNECT MODEL	29

Table of Tables

TABLE 1: GCN INTER-PROVIDER REFERENCE DOCUMENTS.....	8
TABLE 2: PSN TELEPHONY STANDARDS REFERENCE DOCUMENTS.....	27
TABLE 3: TDD – TECHNICAL ABBREVIATIONS	31

1 Document Scope

The Technical Domain Description is initially released by the PSN Programme as a Technical Document that:

- Clarifies the application of standards identified in the Operating model documents
 - Inter-Provider MPLS Network to Network Interface model

- Identifies Standards for Basic Telephony interoperation
 - Simple Telephony Model
 - Interoperation requirements
 - Clarification of standards

- Logical Constrained Resource Management
 - NNI logical addressing
 - DNS
 - Domain Names
 - Numbering
 - Consumer IP Address Management

No information in this document is either confidential or commercially sensitive from the perspective of companies participating in its development. It has therefore been for HMG to determine what confidentiality provisions shall apply to this document. Cabinet Office, the relevant authority in this regard, has determined that this document does not require a protective marking, and the document is marked UNCLASSIFIED to positively indicate this status. Any and all tradenames and trademarks contained in this document are hereby acknowledged as tradenames and trademarks of their registered owners.

1.1.1 Document Relationship to the PSN Operating Model

In its initial release this document is a programme document that supplements the PSN operating model providing necessary clarifications to assist programme stakeholders in interpreting the standards identified in the operating model. This document also supplements the operating model in identifying additional standards for interoperation of PSN Telephony Services. At the time of release a number of agreed changes are pending for the operating model, it is intended that this document will be aligned with the revised operating model and re-released as a component document of the operating model through the applicable document and assurance controls.

2 Management of Constrained Logical Resources in the PSN

Constrained Logical Resources include addressing, names, numbers and labels. There are two primary types of these resources; Public and Private. Public resources are globally unique and are therefore administered by a hierarchy of authorities. Private resources are administered by a private authority which may be an enterprise or a carrier or a nominated organisation. This is necessary for any logical resource that must be unique within the private environment. PSN consumes both Public and Private logical resources. In Private systems there may be mitigations that remove the necessity for uniqueness however such mitigations can add significant cost.

2.1 Public Logical Resources

For the Use of Public resources within the PSN the PSN Authority will be responsible for monitoring the allocation of such resources within PSN. Such allocations may be requested by suppliers or consuming organisations from the appropriate public body, the PSNA or its nominated body provides oversight for PSN logical resources as a whole.

2.2 Private Logical Resources

The PSNA (or its nominated representative) is responsible for the management, coordination and allocation of Private resource allocation within PSN. The key principles the PSNA will apply in discharging this duty are:

- Minimum barrier to entry in supporting transition of existing systems to PSN
- Allocations that trend towards least complexity and cost
- Re-use of existing arrangements and bodies where possible
- Adoption of common industry practises

Logical Resources are consumed by other strands of the Public Sector ICT strategy, for example G-CLOUD and Desktop. The PSN programme will engage with the appropriate cross domain working groups and bodies to ensure that shared logical resources are managed to an agreed practise.

2.3 Instances of Private Logical Resources within the PSN Standards

The following are examples of constrained logical resources that are referenced in the PSN Operating Model and Standards.

- IP addresses
- AS numbers
- Domain names
- SIP addresses
- Telephony numbering plan

2.4 DNS Services and Inter-Impact Level (IL) Domain Browsing

The service provider community offers a wide range of Domain Name System (DNS) mechanisms for use in the resolution of name queries across organisations. Although, these mechanisms comply with a number of industry standards and best practice methodologies a further appreciation of the DNS feature as applied to both current deployments and the intended PSN is required.

The following sections further detail these considerations with regards to the overall logical addressing and environmental specific elements.

2.4.1 Current Public Sector Scenarios

Typically, most GSI connected departments utilise three separate DNS services, providing the following functions:

- Privately addressed Internal DNS
- GSI DNS for all .gsi domains
- Internet-facing DNS resolution (e.g. .gov.uk) with domain names being provided by the applicable Internet Access Supplier, not the GSI.

In addition, current mechanisms for name resolution between different impact levels introduces further complexity and support issues requiring reactive management on a case by case basis. The emergence of the PSN will also amplify the current complexities if not addressed in a systematic manner with appropriate architectures.

As the PSN is required to operate in an 'any-to-any' environment, the requirement for a multi-organisational, multi-provider, cross government DNS and related domain management policy will be needed.

2.4.2 PSN Hierarchical DNS Model

Possible requirements include those for a central DNS authoritative service to facilitate improved security and agility in support of transition and machinery of government changes.

As a result, all government agencies and participating Industry members are being consulted regarding the overall impact of the DNS services and related dependencies in the PSN environment. It is intended that this process will lead to further recommendation on the creation of a unified DNS model aligned with HMG ICT strategy.

2.4.3 Inter Impact Level Domain Recursion

Currently there are separate DNS instances for each IL with no recursion between the layers. Inter IL level gateways then do the lookups for the lower domain if required i.e. IL4-> IL3. In the case of web browsing from IL4 the user would browse at IL4 using the IL4 DNS servers. Should an end user wish to browse down to IL3, the proxy in the IL4 to IL3 gateway would then place a request to the lower domain.

2.4.4 Authoritative DNS Root Servers and the PSN

Currently, Citizen facing domain names are referenced via JANET hosted DNS mechanisms. GSi family based domains are referenced via C&W DNS instances at each impact level and for each customer via existing service provider impact level DNS instances. This is expected to remain unchanged in the PSN environment as it evolves.

2.4.5 Domain Naming Conventions and Explicit Impact Level Identifiers

Current GSi domain naming conventions explicitly identify IL2, IL3 and IL4 domains via the naming structure .gsi.gov.uk, .x.gsi.gov.uk etc.

In order to facilitate our policy of minimum barrier to entry the PSN will support existing domain names wherever possible. However, the PSN will define naming conventions for any new names.

Inter consumer organisation application sharing and communication will increasingly require mechanisms that support the integrity of impact level traffic types including e-mail and media services (i.e. IL3 to IL3, IL4 to IL4 maintained) .

2.4.6 PSN Service Instances Impacted by DNS and Domain Management

There is emerging visibility of a range of services requiring the facility to locate and map inter-service platform registration such as inter SIP server location in stances where numbers of servers are likely to be in the hundreds if not thousands.

3 GCN – Inter Provider MPLS

The initial section of the document focuses on the key deployment models as recommended for the GCN by industry. The following options have been evaluated:

- Inter-AS MPLS Option A
 - Control and Data Plane operations
 - Scaling
 - Security
- Inter-AS MPLS Option B
 - Control and Data Plane operations
 - Scaling
 - Security

Note: Although RFC 4364 defines a number of peering methods, several deployment models have been excluded due to their inability to meet the participating GCN provider's architectural requirements in full. As a result, the Carrier Supporting Carrier (CsC) and Inter-AS options C (Mp-EBGP with Router Reflectors (RR)) architectures have been excluded from the outset.

The second part of the document addresses the concerns around peering with an external carrier with regards to providing the required levels of service transparency and operational controls between GCN service providers. The following elements have been considered:

- GCN Inter-AS provisioning
 - Baseline Architecture
 - Deployment Scenarios

The final part of the document is dedicated to the evaluation of the scenarios against both vendor and carrier best practice methodologies as advocated by Industry.

3.1 Assumptions

It is expected that the target audience for this document fully understands the requirement for the Government Conveyance Network (GCN) to provide peering services and associate Network to Network Interconnects (NNI) to all GCN compliant service providers.

All options for Inter-Provider connectivity correspond to options mentioned in draft RFC 2547bis and the associated RFC 4364 update.

3.2 Reference Documents

The following documents have been referenced in the preparation of GCN Inter-Provider section of the TDD:

Document Ref.	Description	Author
RFC:2547bis (DRAFT)	BGP/MPLS IP VPNs	E. Rosen
RFC:2858	Multiprotocol Extensions for BGP-4	T. Bates
RFC:3270	Multi-Protocol Label Switching (MPLS) Support of Differentiated Services	F. Le Faucheur
RFC:4364	BGP/MPLS IP Virtual Private Networks (VPNs)	E. Rosen
RFC:4381	Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs)	M. Behringer
RFC:4684	Constrained Route Distribution for BGP/MPLS IP VPNs	L. Martini
CESG NGN v9	NGN Standard Security Level – ISMS requirements and Good Practice Guide	CESG
GCN Service Description v1.1	GCN Service Description: Addendum to PSN Operating Model	PSN Core Team

Table 1: GCN Inter-Provider Reference Documents

3.2.1 RFC Notes

The Internet Engineering Task Force (IETF) released the first RFC for MPLS L3VPNs (RFC 2547 – BGP/MPLS VPNs) in 1999 as a draft standard. Further revisions of this draft released in 2003, 2004 and 2005 added additional support for Inter-Domain MPLS L3VPNs functionality. In 2006, RFC 4364 ‘BGP/MPLS IP Virtual Private Networks (VPNs)’ was approved by the IETF. This documented fully updated draft RFC 2547.

RFC 4364 describes three options for Inter-AS or Inter-Provider MPLS VPN solutions - A, B and C. RFC 4381 provides detailed analysis and evaluation of the security issues related to provisioning client VPNs in the BGP/MPLS environment.

4 The Government Conveyance Network (GCN)

The GCN ‘network-of-networks’ model facilitates the interconnection of compliant service providers in order to create a seamless backbone transit architecture that will allow for the provisioning of network access across the UK. This will allow a number of resilient, high-availability Points of Connection (PoC) to be presented for use by any PSN service provider or public sector organization that complies with the associated Code of Interconnection (CoIco). In this way, compliant parties may become Direct Network Service Providers (DNSPs) using the GCN for core transmission and connectivity.

The overall architectural objectives for the GCN are fully detailed in the ‘GCN Service Description v1.1’ document which can be found under the Public Sector Network (PSN) section of the Cabinet Office Website as follows:

http://www.cabinetoffice.gov.uk/cio/public_sector_network/governance.aspx

4.1 GCN - Conceptual Architecture

As described previously, the GCN allows for the provisioning of network connectivity and raw bandwidth services for use by the dNSP community. The following diagram provides a conceptual overview of this architecture:

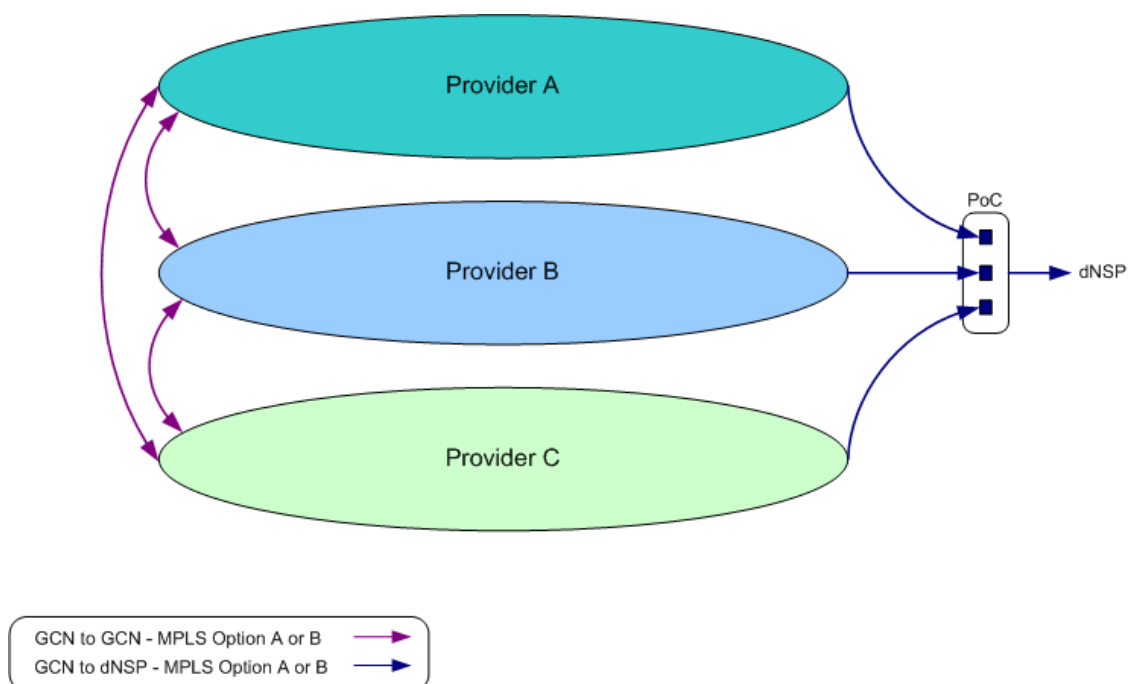


Figure 1: GCN – Conceptual Architecture

Two types of connection exist within the GCN connectivity model.

- GCN to GCN Provider Interconnects:
Required between participating GCN providers for 'full mesh' connectivity to the core.
- GCN to dNSP Interconnects:
Required at the Point of Interconnect for dNSP access to the GCN.

For both types of interconnect Industry has proposed the use of Multi-Protocol Label Switching (MPLS) Inter-Provider connectivity using either 'option A' or 'option B'.

Note: Although GCN service providers may seek to offer either option A or B at the dNSP interconnect, support for both options within the GCN core is mandatory.

The following sections fully evaluate the intended GCN Inter-Provider MPLS Options and associated best-practices as recommended by Industry and validated by the applicable product vendors and service provider community.

5 GCN Inter-Provider MPLS

The GCN 'network-of-networks' model facilitates the interconnection of compliant service providers in order to create a seamless backbone transit architecture that will allow for the provisioning of network access across the UK. This will require all participating service providers to implement peering points within their existing infrastructure(s) for interconnectivity purposes. Naturally, this model introduces several interoperability considerations that need to be fully evaluated prior to any deployment activity.

Implementation issues essentially involve the requirement for all WAN MPLS services to be fully accessible via different autonomous systems located in disparate third-party carrier infrastructures. The Inter-Provider model facilitates this connectivity via mutual redistribution and hand-off of the required protocol parameters in a controlled, secure fashion thereby enabling administratively diverse clients to communicate freely with each other and without issue.

Two options exist when interconnecting independently managed MPLS domains over third-party infrastructures:

- Inter-Autonomous System (Inter-AS) Connectivity
- Carrier Supporting Carrier (CsC) Connectivity

Inter-AS is a peer-to-peer deployment model that allows for the extension of VPNs through multiple provider or multi-domain networks. This solution enables Service Providers to peer with one another offering end-to-end VPN connectivity over extended geographical locations. Inter-AS services adhere to the operational and deployment principles as historically defined in RFC 2547bis and as updated in RFC 4364. This RFC specifies three implementation options for Back-to-back VRF, Multiprotocol eBGP (MP-eBGP) for VPNv4 and Multi-hop MP-eBGP between route-reflectors (RRs) respectively.

Carrier Supporting Carrier (CsC) is a hierarchical VPN model that allows Service Providers, or customer carriers, to directly connect their IP or MPLS networks over an MPLS backbone. This deployment model is typically used by carriers in order to extend their network into areas where they have no regional access due to limited infrastructure or regulatory issues.

Note: RFC 4364 Carrier Supporting Carrier (CsC) and Inter-AS option C (Mp-EBGP with Router Reflectors (RR)) architectures require additional elements (intermediate nodes and IP Addresses) of the provider infrastructure to be exposed between peers. In some scenarios this may pose additional security concerns.

As a result, these deployment models have been deemed non-applicable by Industry and have been excluded from the outset.

5.1 Inter Autonomous System (Inter-AS) MPLS

The following sections detail the key operational and deployment requirements for Inter-AS Option A and B provisioning at the GCN Network to Network Interface (NNI).

5.2 Option A: Back-to-Back VRF

The back-to-back VRF approach is the simplest method for allowing MPLS VPN providers to exchange VPN routing information across different MPLS domains.

VRF instances are configured back-to-back between the ASBR routers and provide direct connectivity between provider Edge (PE) routers in different autonomous systems (ASs). In this respect, the PEs function as ASBRs from WAN to carrier network. ASBRs are connected via dedicated links consisting of logical subinterfaces or multiple physical links. Each physical interface or subinterface is dedicated to a single VRF instance.

The following diagram illustrates the Back-to-Back VRF configuration:

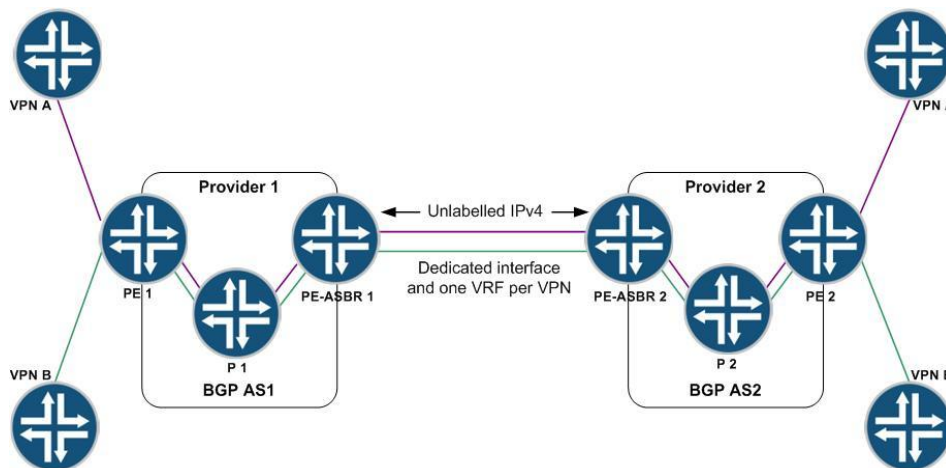


Figure 2: Back-to-Back VRF – Connectivity

The ASBR routers in each autonomous system serve as gateways that are used to exchange IPv4 routes. The link between peer PE-ASBRs will support the external Border Gateway Protocol (eBGP) at the NNI. Typically, eBGP is implemented between ASBRs as this protocol easily scales to the service provider model whilst providing enhanced features such as traffic policy enforcement, scalability, and security mechanisms. VPN routes are sent in IPv4 format.

5.2.1 Packet Forwarding

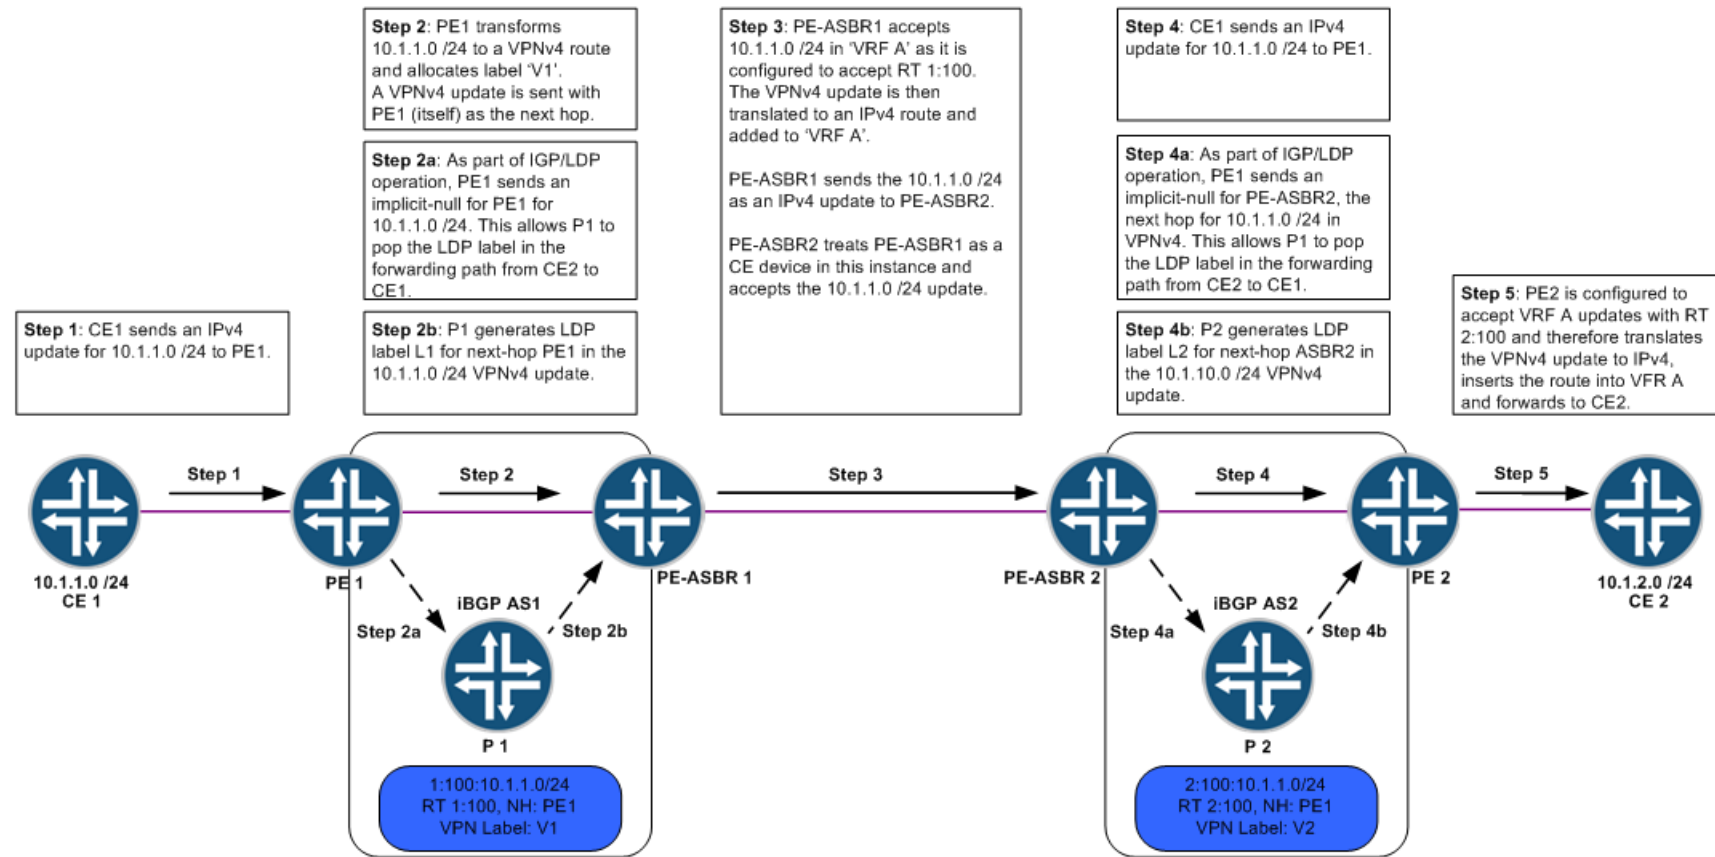
Standard MPLS VPN packet forwarding processes remain unaffected within the provider AS. At the boundary of the network, all Interior Gateway Protocol (IGP) and VPN labels are removed by the associated PE-ASBR. The IPv4 packet is then transmitted to the peer provider PE-ASBR in the adjacent autonomous system.

In effect, the receiving PE-ASBR sees the request as originating from a Customer Edge (CE) router and accepts the packet. MPLS and IP VPN encapsulation is performed and the packet transmitted through the third party carrier backbone AS.

Figure 3: Back-to-Back VRF – Control and Data Plane illustrates both Control and Data plane packet forwarding operations in a Back-to-Back VRF configuration.

Note: Although Control and data plane operations for VPN and IGP label distribution occur independently, both functions are shown in parallel in order to provide full visibility of the end-to-end packet forwarding process.

Control plane



Data plane




Figure 3: Back-to-Back VRF – Control and Data Plane

5.2.2 Security

Within an Option A architecture, each PE-ASBR perceives the other PE-ASBR as a CE. The PE-ASBR does not accept labelled packets from the CE at the ingress point into the network. No signalling is in place between the PE-ASBRs other than inside the VRFs on both sides. Therefore, the separation of the VPNs on both sides and associated security levels are identical to those found in a single AS network.

Since the VPNs are separated on different interfaces, all signalling between ASBRs remains within a given VPN. This means that dynamic cross-VPN security breaches are impossible. Although it is conceivable that a service provider may connect a specific VPN to the wrong interface (allowing access via a non-trusted VPN source), for this to occur, additional configuration steps would be required in order to allow any new host routes to propagate through the network. Inter-Provider connections of this type would typically be secured with filters to limit connectivity to specific network routes and autonomous system paths. In this way, the impact of a simple physical patching error may be mitigated at source.

Essentially, the two parties must trust each other and ensure that adequate operational controls are enforced at PE-ASBR security boundaries.

5.2.3 Scaling

Each VRF instance requires a dedicated interface (physical or subinterface), a VRF per VPN, and a separate IGP route table per VRF be configured on the PE-ASBR routers. As a result, the back-to-back VRF solution can prove resource and configuration intensive as VPN numbers increase. In addition, scaling factors must be considered in order to ensure correct operation of the PE-ASBR routers under load conditions.

5.3 Option B: ASBR-to-ASBR Multi-protocol eBGP for VPNv4

The ASBR-to-ASBR approach uses MP-eBGP to transport VPNv4 prefixes over separate autonomous systems. In this way, the standard routing mechanisms of eBGP can be applied without the need to have dedicated, per-VPN configurations as is the case with the back-to-back VRF configuration detailed in section [5.2 Option A: Back-to-Back VRF](#).

As communication is limited to VPNv4 label and prefix updates, no label operation is needed and therefore, no IGP is required between ASBRs. All VPNv4 labelled packets are forwarded by the MP-eBGP session.

The only difference between MP-iBGP and MP-eBGP when transporting VPNv4 prefixes is the way the next-hop attribute is handled. As this value changes at the ASBR, the LSP path is terminated and a new VPNv4 label added. The VPNv4 label is then sent as an MP-eBGP update to the neighbouring ASBR peer.

The following diagram illustrates the ASBR-to-ASBR MP-eBGP configuration:

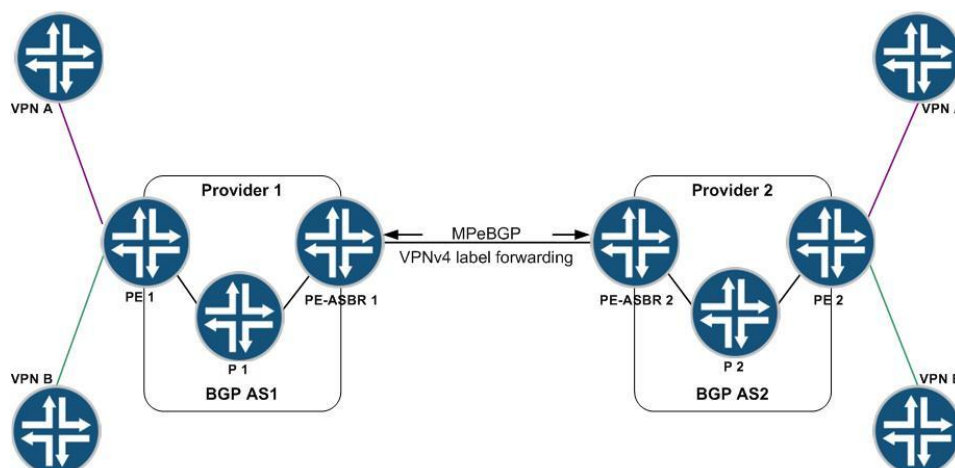


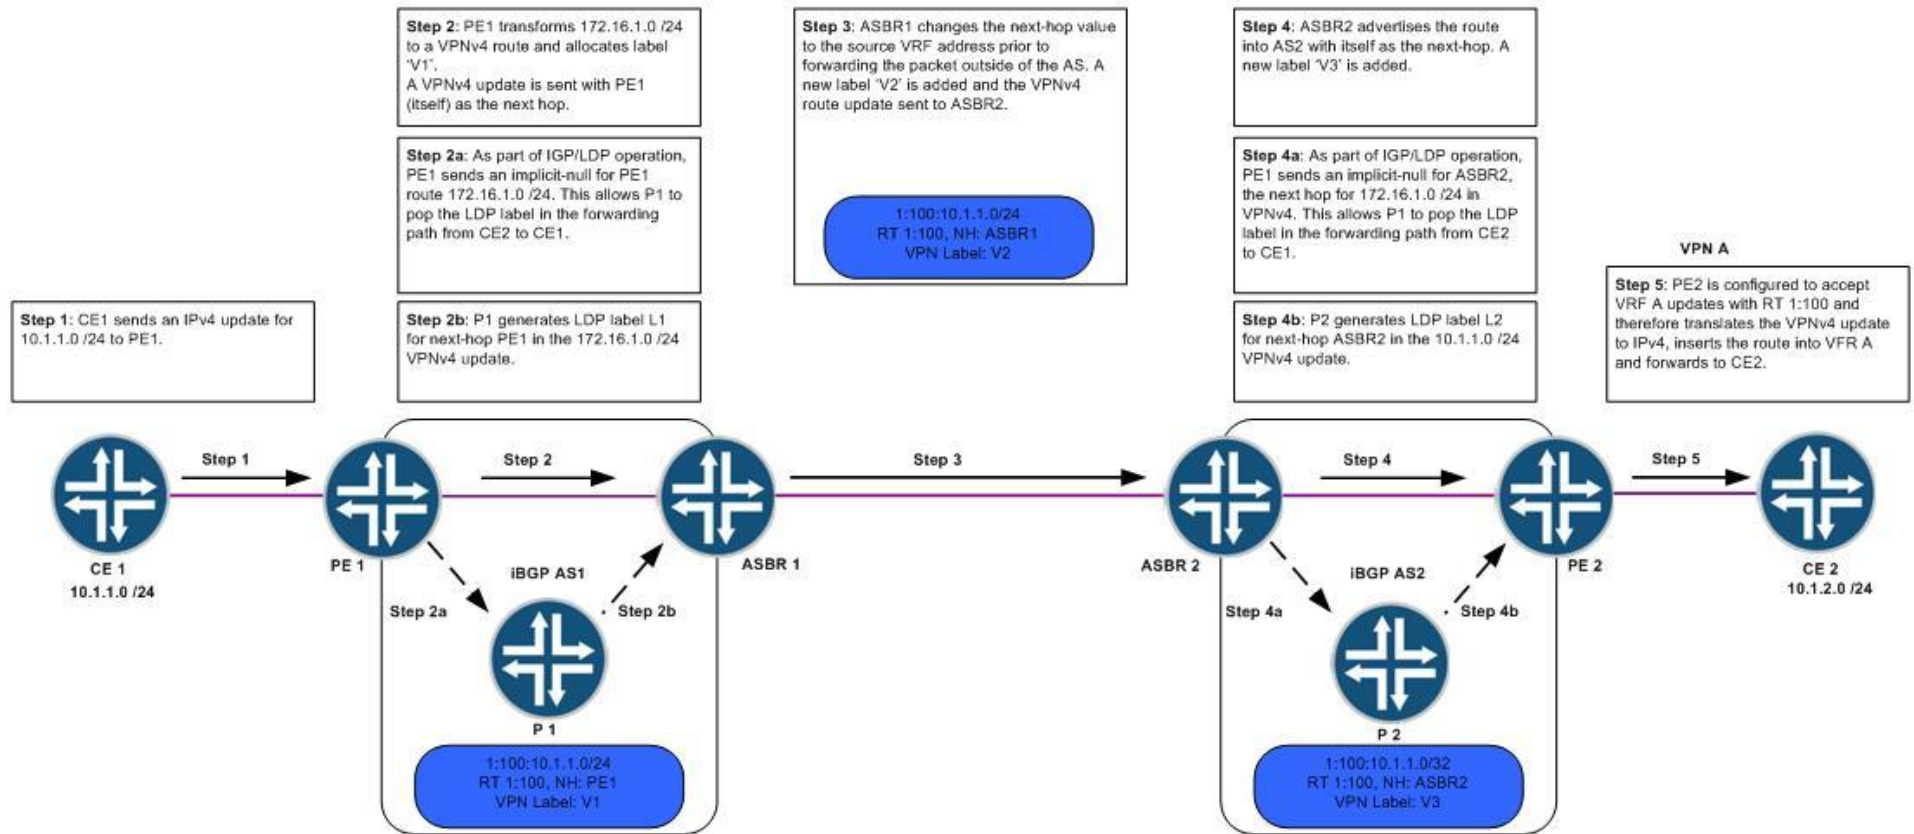
Figure 4: ASBR-to-ASBR – Connectivity

5.3.1 Packet Forwarding

Standard MPLS VPN packet forwarding processes remain unaffected within the provider AS. At the boundary of the autonomous system, the next-hop field is modified and a new VPNv4 label applied. MP-eBGP then forwards the labelled packet onto the neighbouring ASBR who in turn changes the next-hop attribute to itself prior to forwarding the labelled packet (and new VPNv4 label) into the MP-iBGP autonomous system.

Figure 5: ASBR-to-ASBR – Control and Data Plane illustrates both Control and Data plane packet forwarding operations in an ASBR-to-ASBR next-hop-self configuration.

Control plane



Data plane

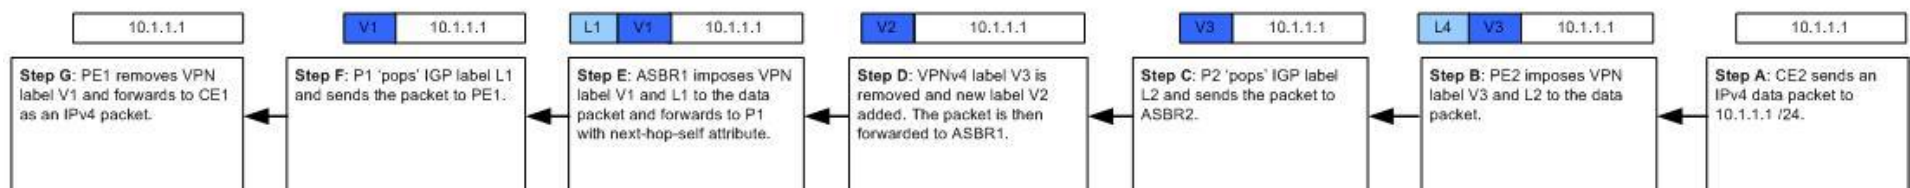


Figure 5: ASBR-to-ASBR – Control and Data Plane

5.3.2 Security

Option B deployments, despite being scalable, do not provide perfect VRF isolation as all configuration is limited to a single interface when forwarding labelled packets. All eBGP routing information for the configured VPNs is made available to peer ASBRs in this scenario. This is not held in separate VRFs, but in the MP-eBGP database itself. In all cases, this should be limited to the Inter-AS VPNs through filtering. Logical separation within the PE is maintained through the use of VPNv4 addresses and is subject to standard MPLS VPN considerations (synonymous with Layer 2 circuit technologies).

The control plane between the ASBRs uses Multi-Protocol eBGP. VPN routes are exchanged as VPNv4 addresses with the ASBR in each autonomous system acting as the next-hop for IPv4 peering. The data plane is separated through the use of a single label, representing a VRF or a subset thereof.

RFC 4364 states that an ASBR should only accept packets with a label that it has assigned to this router (the ASBR). This prevents the insertion of packets with unknown labels, but it is possible for a service provider to use any label that the ASBR of the other provider has passed on. This may allow one provider to insert packets into a provider VPN for which it has a label.

This solution also needs to consider security at the layer 2 interconnection point. Therefore, this type of interconnection will only be implemented at private interconnection points. This will be achieved via a fully secured private point-to-point connection that cannot be interfered with by any third party.

Essentially, the two parties must trust each other and ensure that adequate operational controls are enforced at the ASBR security boundaries.

5.3.3 Scaling

Physical scalability with this model is greatly improved over the Back-to-Back VRF model as dedicated interface (physical or subinterface), a VRF per VPN and a separate IGP route table per VRF is no longer required.

In this scenario, a single eBGP session is all that is required between ASBRs however, the ASBR still needs to maintain a full routing table for all Inter-AS VPNs. This may become a scalability factor in larger Inter-AS deployments. Scaling factors should be tested in order to ensure correct operation of the ASBR routers under expected load conditions.

5.4 Quality of Service (QoS)

RFC 3270 'Multi-Protocol Label Switching (MPLS) Support of Differentiated Services' defines a solution that allows service providers to map Differentiated Services (Diff-Serv) Behaviour Aggregates (BA) onto MPLS Label Switch Paths. Both Inter-AS deployment Options A and B are all fully RFC 3270 compliant.

In this model, the IP Precedence field (the first three bits of the DSCP field in the header of an IP packet) specifies the QoS. Based on this marking, the packet is given the desired treatment such as the latency or the percent of bandwidth allowed for that class of service. If the service provider network is an MPLS network, then the IP Precedence bits are copied into the MPLS EXP field at the edge of the network thus preserving QoS values across the provider network.

If the service provider wants to set the QoS of an MPLS packet to a different value than that of the IP Precedence bits, the service provider can set the MPLS EXP field instead of overwriting the value in the customer's IP Precedence field. The IP header remains available for the customer's use and is not changed as the packet travels through the MPLS network.

Although the GCN provider may support a number of separate mechanisms in order to provision traffic prioritisation at the network edge, as may be the case with interfaces to direct Network Service Providers (dNSPs), all traffic types will be grouped by their respective handling requirements. For example, expedited

forwarding (EF) traffic may appear as part of a 5-class carrier policy or as a six-class client policy however, these parameters can be mapped to the correct class at the provider edge without impact.

Further dialogue with the carrier will be required during the planning phase to confirm the bandwidth ratios, queues and subsequent QoS mapping parameters.

6 GCN Inter-AS Deployment Model

The previous sections detail the underlying architectural standards associated to Inter-As Option A and B solutions. These standards generally define the interconnect requirements in a single-hop point-to-point environment. As a result, an appreciation of the underlying architecture as proposed for use in the GCN infrastructure is required.

The following section details the application of these standards within the intended GCN deployment model.

6.1 Baseline Architecture

Although all GCN service providers are required to conform to standards based deployment at the Network to Network Interface (NNI) point, the internal service provider architectures themselves are considered proprietary and out-of-scope for evaluation purposes. However, in order to assure end-to-end support of the MPLS Inter-AS feature in a multi-hop environment, a representative Baseline Architecture must be assumed.

The following diagram illustrates the Baseline Architecture as presented and agreed in the GCN work stream. In this way, it is possible to consider the impact of any NNI feature(s) on the overall data path end-to-end without compromising the intellectual property of any participating service provider.

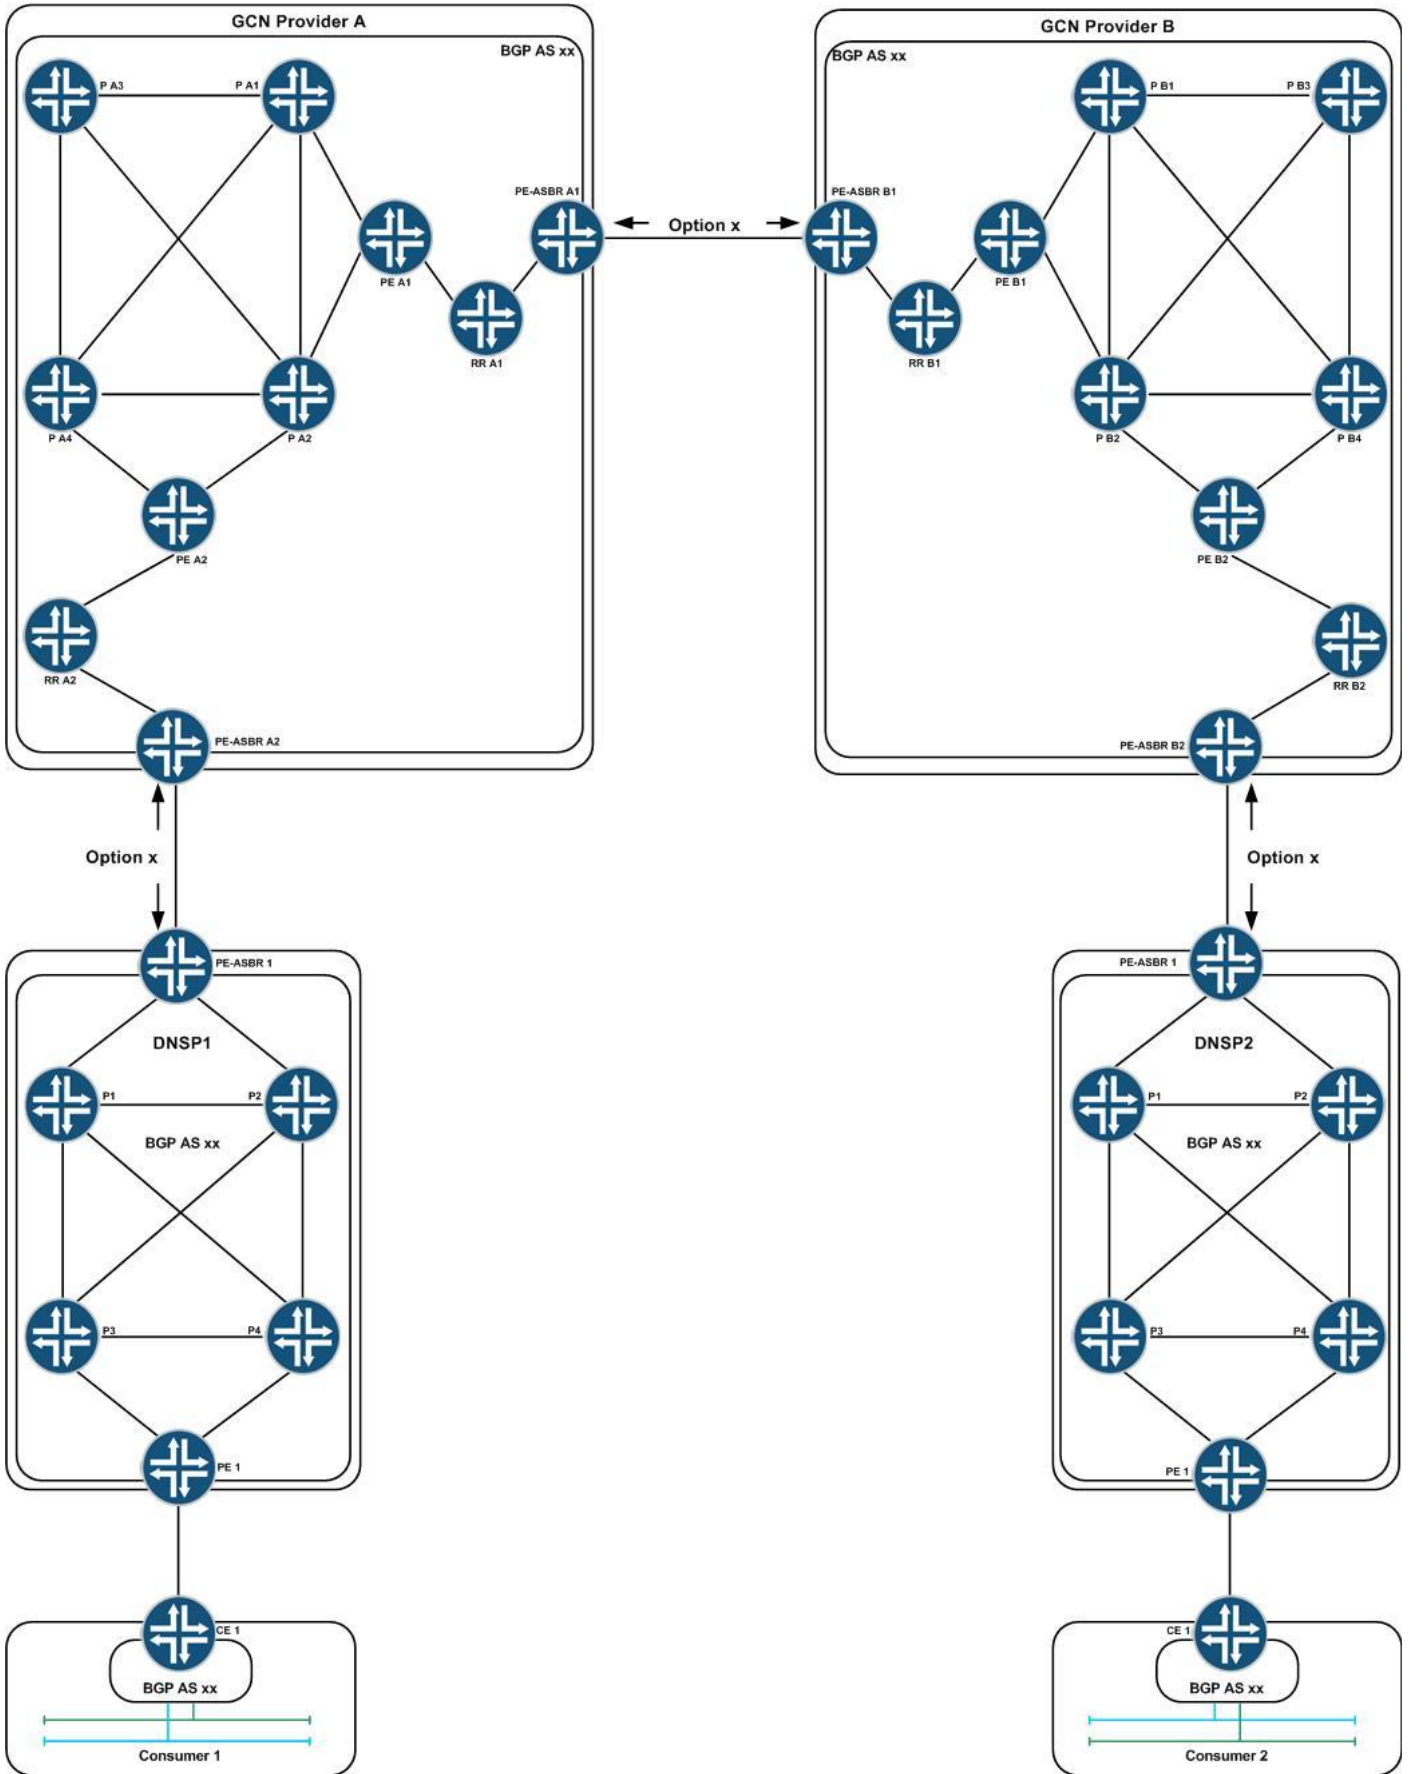


Figure 6: GCN – Generic Baseline Architecture

The architecture assumes a minimum level of GCN provider equipment to support the use of three Inter-AS interconnects over the following bi-directional communication path:

Consumer1 > DNSP1 > GCN Provider A > GCN Provider B > DNSP2 > Consumer2

Five deployment scenarios have been evaluated in order to verify correct operation of cross-provider MPLS VPN services using Option A and B interconnects as specified in RFC 4364.

Testing of these scenarios was carried out by a number of independent vendors, directly responsible for the definition of the standards being referenced.

6.1.1.1 Scenario 1: A/A/A

In this scenario, GCN service providers have deployed MPLS Inter-AS Option A at all NNI points in the environment.

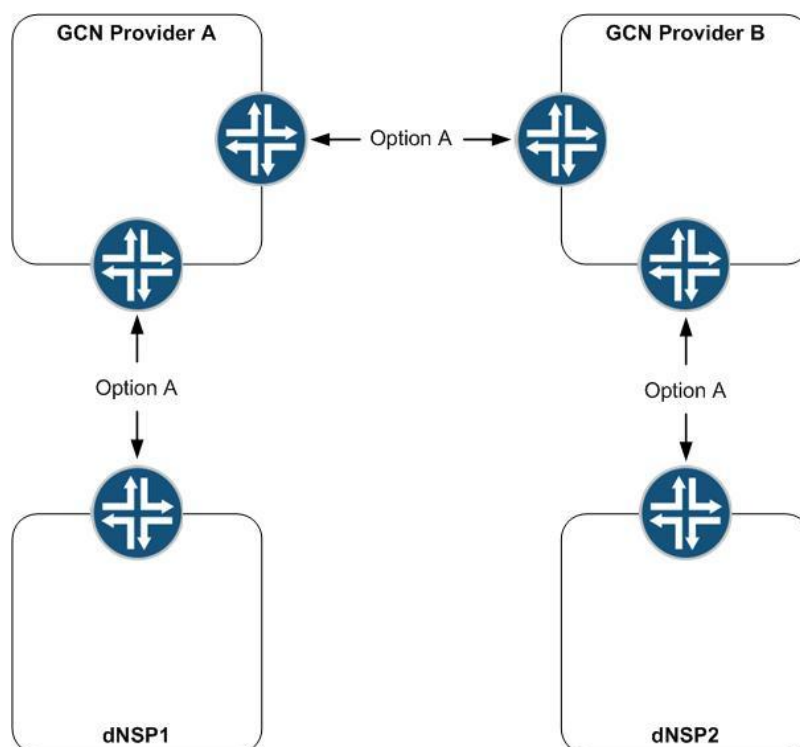


Figure 7: Scenario 1 – MPLS Option A/A/A

Bi-directional traffic flow has been shown to operate without error in this configuration.

6.1.1.2 Scenario 2: A/B/A

In this scenario, MPLS Inter-AS Option A has been deployed at the dNSP NNI points with Inter-AS Option B between GCN service providers.




Figure 8: Scenario 2 – MPLS Option A/B/A

Bi-directional traffic flow has been shown to operate without error in this configuration.

6.1.1.3 Scenario 3: B/B/B

In this scenario, the GCN service providers deployed MPLS Inter-AS Option B at all NNI points.

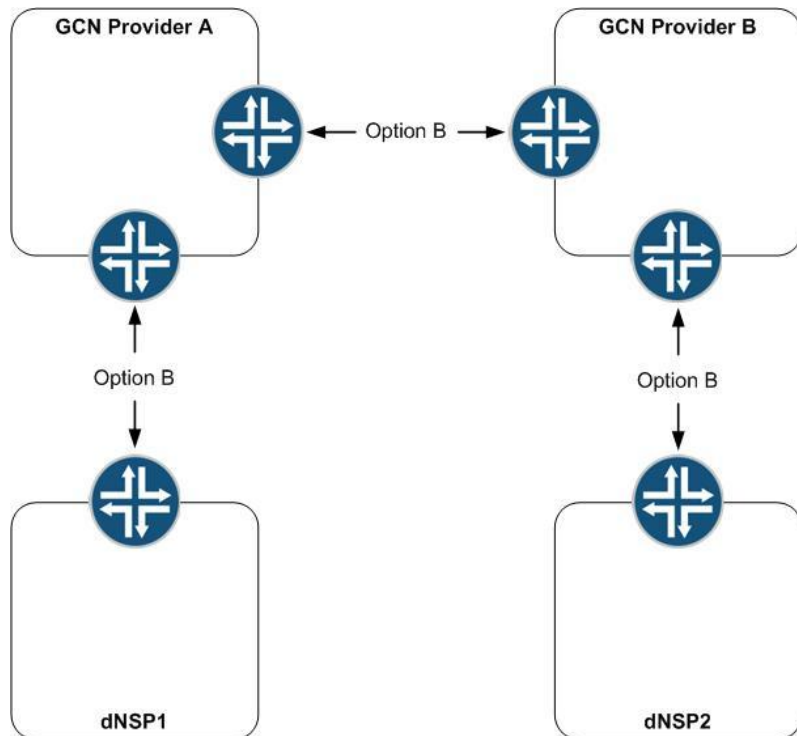


Figure 9: Scenario 3 – MPLS Option B/B/B

Bi-directional traffic flow has been shown to operate without error in this configuration.

6.1.1.4 Scenario 4: B/A/B

In this scenario, MPLS Inter-AS Option B has been deployed at the dNSP NNI with Inter-AS Option A between GCN service providers.

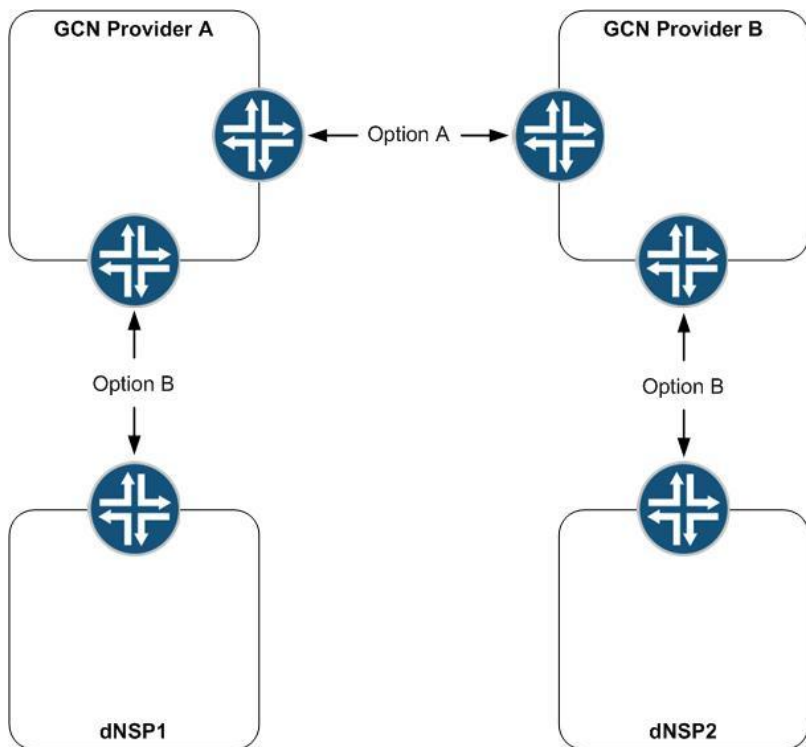


Figure 10: Scenario 4 – MPLS Option B/A/B

Bi-directional traffic flow has been shown to operate without error in this configuration.

6.1.1.5 Scenario 5: Both A/B Interconnect Points

In this scenario, both Inter-AS MPLS Options A and B are present at all interconnecting points. Although the GCN service provider may wish to provide separate infrastructure for these interfaces, this architecture accurately reflects a ‘migration’ state between Option A and B NNI points which may be likely in the event of any Option A scaling issues.

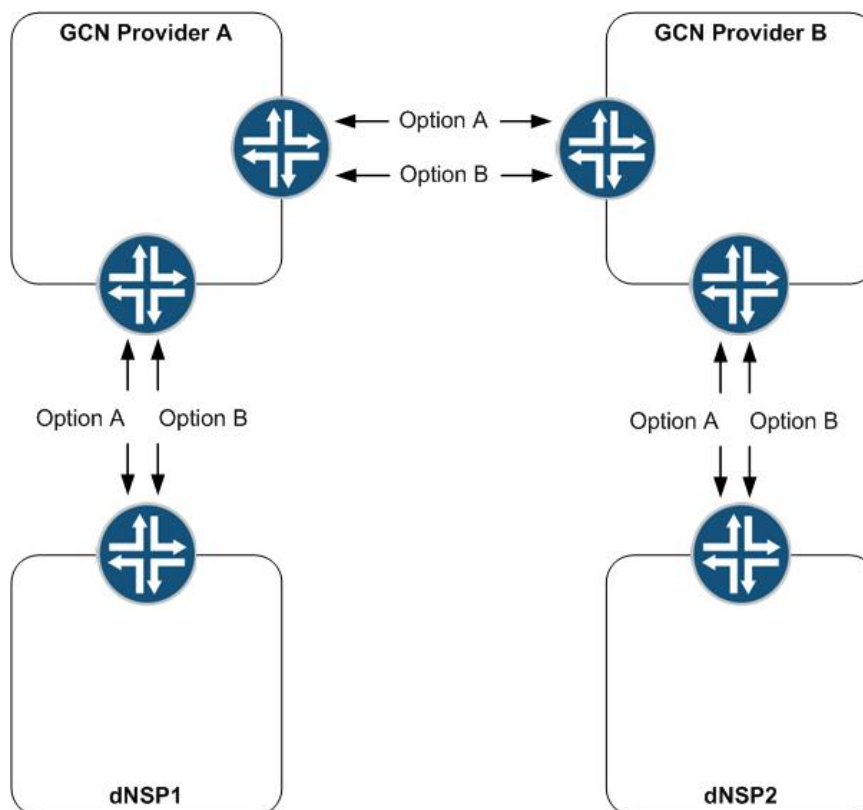


Figure 11: Scenario 5 – Both MPLS Option A/B Interconnects

Bi-directional traffic flow has been shown to operate without error in this configuration however, additional filtering techniques must be deployed in order to ensure the required level of traffic separation.

Note: As this scenario demonstrates full dual-state operation of both options A and B at the NNI, the evaluation of scenarios A/A/B and B/B/A is not required.

7 GCN Network to Network Interface (NNI) - Common Attributes

All GCN service providers agree compliance with the Inter-AS option A and B standards ([3.2.1 RFC Notes – 4364 mandatory](#)) mode of operation (section [5 GCN Inter-Provider MPLS](#)), deployment model (section [6 GCN Inter-AS Deployment Model](#)) and appropriate CESG NGN (2-2-4) assurance requirements in their entirety.

Industry has also recommended a number of additional configuration standards that should apply to the provisioning NNI interconnects in the GCN environment. Incorporation of both elements (standards and recommendations) will allow for a formalized approach to be taken in the deployment of any MPLS Options A and B interconnects to a common technical framework.

The following sections present the implementation specifics and common attributes that need to be applied to any GCN NNI in all cases.

7.1 Physical Deployment

7.1.1 Device Location and Security

All devices forming part of the NNI interconnect will be housed at the respective GCN and dNSP service provider premises unless agreed co-location services are sought.

Irrespective of location, all Service Provider facilities will be physically secured with access only being granted to approved and suitably screened staff.

Equipment housed at these facilities will be suitably secured and isolated so as not to be accessible without strict permissions. Any third-parties or contractors requiring access to the facility will be subject to prior vetting and only permitted access to the required equipment for specific works at agreed times.

7.1.2 Interface Allocation

The GCN service provider will provide a mechanism for the allocation of interfaces at the NNI interconnect. This will allow for the correct level of audit and capacity procedures to be followed. In the event of multiple connections being presented at the NNI, these will be allocated in a redundant fashion where required. This will include physical and/or logical separation where necessary to meet the required provisioning task.

7.1.3 Cabling

The GCN Provider will provision all cabling to the associated termination point. From this point, cabling is expected to be provided via direct patching between peers in a point-to-point fashion. This would not require any further devices, such as switches, to be placed in-line between connections. However, the service provider may be required to add additional infrastructure elements should additional monitoring (taps, probes etc.) be requested. In this event, the service provider will ensure that no performance degradation is incurred as a result.

7.1.4 Labelling

The GCN service provider will ensure adequate labelling of all NNI devices to current standards so that equipment may be easily identifiable in the event of any issues.

7.1.5 Link Failure

All NNI Interconnects will employ monitoring mechanisms to ensure that any link failure is reported in a timely fashion. This failure will be immediately reported into the relevant management platform with an associated alert condition.

7.1.6 Demarcation

In all cases, the GCN service provider control the entire infrastructure to actual cable/fibre or patch-panel hand-off at the dNSP or third-party GCN service provider location.

7.2 Layer 2

7.2.1 Virtual LAN (VLAN) Allocation

VLANs will be allocated from the existing GCN service provider pool. Although maximum VLAN numbers for specific telecoms equipment models are stipulated at a vendor level, these are not considered to represent any limiting factor in deployment of NNI interconnects in the GCN.

7.3 Layer 3

7.3.1 IP Addressing

In all case the GCN service provider is expected to provision NNI interconnects using point-to-point /30 IP addressing presented on physical or logical sub-interfaces. The addition of any further infrastructure as specified in section [7.1.3 Cabling](#) should not affect this numbering scheme.

7.4 Routing Protocols

All GCN service providers have agreed on the use of BGP at the NNI interconnect. The following sections detail the common feature that will apply all BGP session parameters.

7.4.1 Peering

Physical or sub-interfaces will be used for peering in all cases.

7.4.2 AS Numbering and AS Path

GCN service providers will only use public AS numbers. Any AS path manipulation (such as AS-Override, SoO and AS prepend values) is expected to be performed at the edge of the network via the client facing dNSP PE device. No other manipulation of the AS Path is expected.

7.4.2.1 Support for 32 bit AS Number Plan

From January 2011, all newly allocated AS numbers will conform to the 32 bit numbering standard. Although the expectation is for a continuation of the current 16 bit numbering scheme in the short term, all participating GCN service providers have committed to future support of the new numbering standard.

7.4.3 Timers

All GCN service providers support the standard BGP timers. Any enhancement to these values will be agreed between parties for use at the NNI.

7.4.4 Authentication

Where agreed by both parties the GCN service provider will provide authentication of the BGP session at the NNI.

7.4.5 Prefix Limits

In order to maintain NNI resource availability and provide a sub-set of protection against attack at the NNI, all GCN service providers will apply maximum values to the allowed BGP prefix count.

7.4.6 Attribute Support

All GCN service providers will support community values at the NNI point.

7.4.7 Filtering

The GCN service provider commits to applying the correct filters to limit communication to the required VPN resource only. For example, PE devices will allow direct ICMP and BGP peering traffic only.

7.4.8 Load Balancing

All GCN service providers support BGP load balancing mechanisms at the NNI interconnect (MED, BGP-Multipath and AS prepend etc.) dependant on the associated deployment criteria.

7.4.9 VPN Configuration - Option A

All Inter-AS Option A interconnect traffic is passed between peers as native IPv4 packets, therefore, there is no requirement for peer VPN values to match in this scenario. However, all service providers must provide a mechanism for the easy identification and isolation of the required NNI elements across parties in all cases to support co-operation of any troubleshooting activities.

The following parameters will be allocated using the current OSS processes available to the GCN service provider:

- VPN Name
- VPN Description
- VPN Route Distinguisher (RD)
- VPN Route Target (RT)

7.4.10 VPN Configuration - Option B

Inter-AS Option B requires RT and RD values to match at the interconnect. In this case, the GCN service providers will need to agree on the applicable values. In the event that this is not possible the GCN service provider must support the use of the RT re-write feature. All other parameters may allocated using the current OSS processes available to the GCN service provider. However, these systems must allow for the easy identification and isolation of the required NNI elements across parties in all cases to support co-operation of any troubleshooting activities.

7.4.10.1 RD Value

In the event that duplicate RD values are present at the ingress into the GCN (as provided by dNSP or customer networks) the GCN service provider must support RD re-write functionality. This is also required if agreement cannot be reached in the provisioning of RD values at the Inter-AS option B interconnect.

7.4.10.2 RT Re-write

In the event that duplicate RT values are present at the ingress into the GCN (as provided by dNSP or customer networks) the GCN service provider must support RT re-write functionality. This is also required if agreement cannot be reached in the provisioning of RT values at the Inter-AS option B interconnect.

7.4.11 VPN Option A and B Support

In the event that the GCN service provider is required to migrate between Option A and B deployments, it is agreed that this operation will form part of an agreed migration strategy without impact on the overall service.

7.4.12 Scaling

All GCN service providers will apply suitable mechanisms to ensure that the GCN environment and associated components are adequately dimensioned to support the expected (current and forecast) levels of traffic.

7.4.13 Attack Mitigation

The GCN service provider will implement the appropriate controls and best practice methodologies required to mitigate against attack at the NNI point.

7.4.14 Monitoring

All devices at the NNI will support operational reporting mechanism such as SNMP traps and syslog messaging. These will be tuned as necessary to poll for specific elements in both MPLS Option A and B deployments so that the required level of service granularity may be achieved.

In the event of any error condition, an alert will be forwarded to the applicable Network Operation Centre (NOC) toolset along with an associated alarm level. This will trigger the appropriate action to be taken against any pre-defined Service Level Agreements (SLA).

8 PSN Telephony Standards

PSN requires a minimum set of standards that ensure PSN compliant Voice over IP based telephony systems can make calls between them such that the calls stay within PSN. These standards are supplementary to the Security and Connectivity Standards defined in other PSN documents. These standards do not apply to Time Division Multiplex (TDM) Telephony or to voice over IP (VoIP) capabilities of systems that do not have the capability to connect to the Public Service Telephony Network (PSTN), for example; Instant Message clients with voice Over IP capability. Not included in the Telephony standards herein but the subject of future investigations will be the possible need for standards on the usage,

placement and referencing of gateways (Media Gateways and Session Border gateways). The interoperability standards for basic Telephony are described for operating in the clear on PSN, Impact Level 2-2-4. The application of security and encryption policies for higher impact levels will be overlaid on the basic interoperability model and is the subject of future recommendations from the PSN security work stream.




Figure 12: PSN Telephony – multiple departments using multiple services

8.1.1 PSN Telephony Standards - Reference Documents

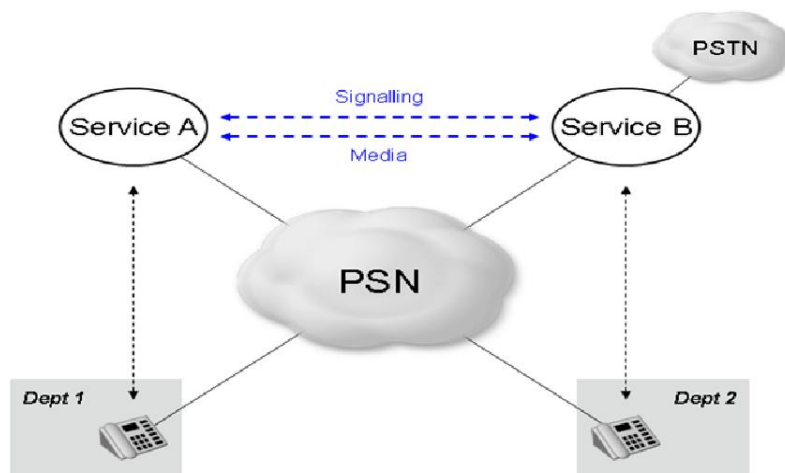
The following documents have been referenced in the preparation of PSN Telephony Standards section of the TDD:

Document Ref.	Description	Author
RFC:3261	SIP: Session Initiation Protocol with refinements	J. Rosenberg
RFC:3264	An Offer/Answer Model with the Session Description Protocol (SDP	J. Rosenberg
RFC:2833	RTP Payload for DTMF Digits, Telephony Tones and Telephony signals	H. Schulzrinne
RFC:3362	Real-Time Facsimile (T.38) – image/t38 MIME Sub-type	G. Parsons
RFC:3325	Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks	C. Jennings
RFC:3323	A Privacy Mechanism for the Session Initiation Protocol (SIP)	J. Peterson
RFC:3550	RTP: A Transport Protocol for Real-Time Applications	H. Schulzrinne

Table 2: PSN Telephony Standards Reference Documents

8.2 Application of Standards

The standards described within this section are for a minimum set supporting Basic Voice over IP connectivity. This is the minimum requirement for Voice Over IP with the purpose of ensuring that calls between parties within the PSN remain on network. It is intended that through procurements, or industry offers, that a range of additional interfaces and capabilities will be specified. These will be according to the client need and would be in addition to the mandatory PSN standards. Although not defined within the Telephony standards, the use of additional interfaces and standards to enhance the capabilities between services within PSN is subject to the general requirements for connection to PSN. Figure 12 shows the general case for Telephony systems interconnect in a PSN context.



Calls between telephony services traverse transparently over the PSN

Figure 13: PSN Telephony – Provider Interconnect

8.3 Use of a Mediation Service

To facilitate transition of existing systems, such as Voice over IP and TDM, PSN will accept as “Telephony Compliant”, a system that achieves the mandatory interfaces by subscribing to a mediation service which converts between the offered interconnect standards and the PSN minimum standard. The mediation service offered in combination with the non-compliant system must be PSN compliant, that is, have the appearance of a compliant system. The endorsement of the system lacking the requisite PSN standard interfaces will be conditional on the continued availability and subscription to the mediation service.

Figure 13 shows the use of a Mediation Service as

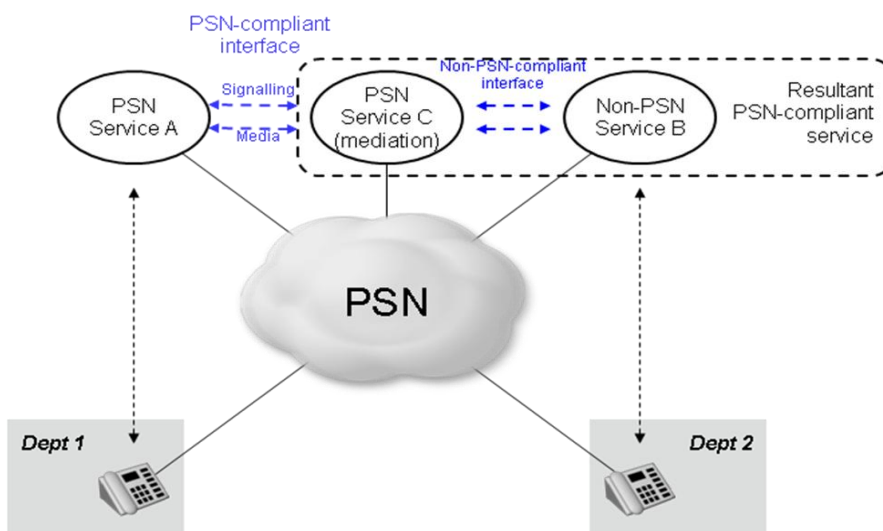


Figure 14: PSN Telephony – Mediation Service

8.4 Interconnect Model

Figure 14 shows the interconnect architecture model for Service provider VoIP peering. This is a logical representation that does not imply a particular implementation within service provider (SP) networks. [Derived from NIC ND1612 functional architecture for PSTN / ISDN generic connectivity (simplified)]

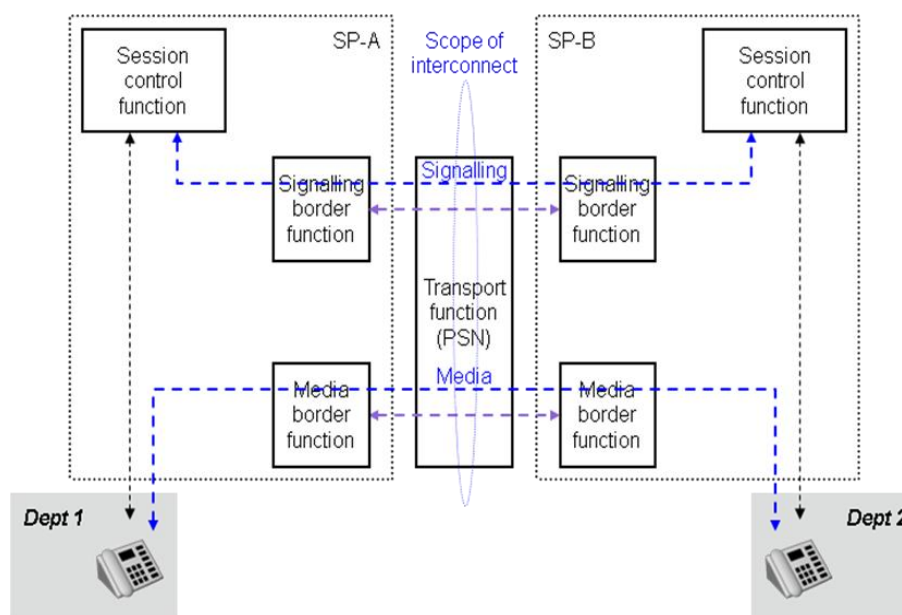


Figure 15: PSN Telephony – Logical Interconnect Model

8.5 Basic Call

For the purposes of this document a basic call is defined as:

- Make a Call
- Receive a call
- Hold a Call
- Transfer a call

8.6 Interoperation Standards

Telephony Servers must support interoperation using RFC 3261: 'SIP Session Initiation Protocol' with refinements as identified in the following sub sections.

8.6.1 Location Identifier

Not specified for PSN compliance but may be requested through specific system specification. PSN may include use of this parameter in the future.

8.6.2 CODEC Selection

Codec selection will be in accordance with RFC 3264: 'An Offer/Answer Model with the Session Description Protocol (SDP)'. To ensure that there is at least one mutually acceptable codec the default servers must support Codec G711 A-Law with 10mS packetisation. To facilitate interoperation existing systems that are constrained to lower bandwidth usage the codec specification G729 must also be supported. Systems may negotiate to establish alternative codec selection but all endorsed systems must support the mandatory codec's.

8.6.3 Negotiation

Negotiation between PSN servers shall be as specified in RFC 3264: 'An Offer/Answer Model with the Session Description Protocol (SDP)'.

8.6.4 Transport Protocol

The Transport Protocol for SIP between PSN servers shall be as specified in RFC 3261 'Session Initiation Protocol'.

8.6.5 Network Authentication

PSN is a single network domain and since all networks within PSN are accredited network authentication is achieved through the use of trusted IP address.

8.6.6 Real Time Protocol (RTP) Transport Media

The media transport protocol for PSN telephony is as specified in RFC 3550 RTP: A Transport Protocol for Real-Time Applications.

8.6.7 Voice Band Services

PSN requires Support for Facsimile and this must be in accordance with RFC 3362: Real-Time Facsimile (T.38). PSN requires that RFC 2833: RTP Payload for Dual Tone Multi frequency (DTMF) Digits, Telephony Tones and Telephony signals is supported. Data Modem support is not required.

8.6.8 Calling Line Identity (CLI)

PSN will support CLI as a default requirement. CLI may be suppressed at the client's request. Transport of CLI will be in accordance with RFC 3325 Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks. (Note: implies RFC 3323 A Privacy Mechanism for the Session Initiation Protocol (SIP)).

8.6.9 Numbering Plan

PSN servers will support fixed and variable number length. It is the intent that systems transitioning to PSN will maintain the external Direct Dial Inward (DDI) numbering geographic or otherwise. The intent of PSN would naturally lead to new numbers being defined as non-geographic however; the future numbering plan for PSN will be the subject of cross government and industry consultation. To facilitate on PSN dialling PSN will define and maintain a numbering plan which uniquely identifies the organisation or group within PSN which is preceded by digit indicating that the call is within PSN (consideration will be given to adoption of the format currently defined for the Government Telephony Network GTN). Refinement of numbering for PSN will be the subject of cross government discussion however; a future numbering plan for PSN telephony should support the formats:

- E164 for public listed numbers
- A Uniform Resource Identifier (URI) for SIP addressing comprising *E164@psn.gov.uk* or similar
- A URI for SIP addressing comprising *name@psn.gov.uk* or similar
- An intra domain and inter domain number/URI comprising a number representing the department or organisation preceded by a unique prefix denoting that the route is within PSN.

9 APPENDIX A – Technical Abbreviations

Term	Description
TDD	Technical Domain Description
PSN	Public Sector Network
PSNA	Public Sector Network Authority
GCN	Government Conveyance Network
DNSP	Direct Network Service Provider
POC	Point of Connection
COICO	Code of Interconnection
GSi	Government Secure Intranet
JANET	Joint Academic Network
RFC	Request For Comment
IETF	Internet Engineering Task Force
IL	Impact Level
IPv4	Internet Protocol Version 4
IGP	Interior Gateway Protocol
BGP	Border Gateway Protocol
MP-eBGP	Multi-Protocol external BGP
MP-iBGP	Multi-Protocol internal BGP
AS	Autonomous System
SOO	Site of Origin
ASBR	Autonomous System Boundary Router
RR	Route Reflector
MPLS	Multi-Protocol Label Switching
LSP	Label Switch Path
CE	Customer Edge - Network Device
PE	Provider Edge - Network Device
P	Provider - Network Device
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
RD	Route Descriptor
RT	Route Target
CSC	Carrier Supporting Carrier
NNI	Network to Network Interface
QoS	Quality of Service
Diff-Serv	Differentiated Services
DSCP	Differentiated Services Code Point
BA	Behaviour Aggregate
MPLS EXP Field	Experimental Field
EF	Expedited Forwarding
VLAN	Virtual Local Area Network
WAN	Wide Area Network
LAN	Local Area Network
NOC	Network Operations Centre
SNMP	Simple Network Management Protocol
TDM	Time Division Multiplexing
PSTN	Public Service Telephone Network
VOIP	Voice Over IP
SIP	Session Initiated Protocol
SDP	Session Description Protocol
DTMF	Dual Tone Multi Frequency
RTP	Real-Time Protocol
MIME	Multipurpose Internet Mail Extensions
CODEC	Code/Decode
CLI	Calling Line Identity
DDI	Direct Dial Inward
URI	Uniform Resource Identifier

Table 3: TDD – Technical Abbreviations