

JOINT OPERATORS TECHNICAL SPECIFICATION OF THE NEUTRAL HOST IN-BUILDING SMALL CELL SOLUTION

ANNEX 1 ARCHITECTURE

SCOPE

This Annex defines the Architecture specification for the JOTS Neutral Host In-Building (NHIB) solution, capable of supporting cellular services for multiple Mobile Network Operators.

This Annex sets out the roles and responsibilities across three operational domains within the Neutral Host In-Building solution. Specifically, the **Operator Domain**, the **Neutral Host Domain** and the **Retailer Domain**.

This Annex describes the security model mandated for the Neutral Host In-Building solution.

PURPOSE

This specification will be used by *Operators*, *Neutral Hosts* and *Retailers* to implement instances of the Neutral Host In-Building solution. To assist in that task the overall specification is divided into a set of annexes, each covering a key aspect of the implementation:

- Annex 1 – Architecture (**This document**)
- Annex 2 – Radio Requirements
- Annex 3 – Testing and Acceptance
- Annex 4 – Operational Processes
- Annex 5 – Fulfilment

Each annex is separately version controlled. Collectively the latest versions of all the annexes define the JOTS Neutral Host In-Building specification.



JOTS 
(NEUTRAL HOST IN-BUILDING)

ANNEX 1
ARCHITECTURE

ALL RIGHTS RESERVED

This is an unpublished work. No part of this document may be copied, photocopied, reproduced, translated, or reduced to any electronic or machine-readable form without the prior permission of the JOTS forum.

DOCUMENT INFORMATION

Document Name:	JOTS NHIB Specification Annex 1 - Architecture
Brief Description:	JOTS NHIB Specification
Document Author:	Graham Ayre (Virgin Media O2)
Owner While Current:	David Morris (Virgin Media O2)
Owner's Email Address:	david.morris@virginmediao2.co.uk
Issue Date:	1 st June 2022
Document contributor	David Morris (Virgin Media O2)
Document contributor	Ada Lam (BTEE)
Document contributor	Doug Fripp (BTEE)
Document contributor	Jon Van-Orden (NEC)
Document contributor	Ranjith Palanivelu (NEC)
Document contributor	Martin Kirk (BT Wholesale)

ACKNOWLEDGEMENT

This document is created with inputs and contributions from the current UK mobile network operators Virgin Media O2, Vodafone, BT/EE and Three.

TABLE OF CONTENTS

1	INTRODUCTION	10
2	DOMAINS	11
3	CONNECTIVITY	15
3.1	VENUE CONNECTIVITY	15
3.2	VENUE TYPES	17
3.3	F-INTERFACE CONNECTIVITY	20
3.4	B-INTERFACE CONNECTIVITY	22
4	ROUTING	24
4.1	ROUTING OPTIONS	24
4.2	EBGP PEERING	28
4.3	B-INTERFACE ROUTING AND RESILIENCE	30
4.4	IP ADDRESSING	34
4.5	DNS/FQDN	38
5	TRAFFIC AGGREGATION	40
6	DISAGGREGATED RAN	42
7	SECURITY	45
7.1	SECURITY DOMAINS	45
7.2	IP SECURITY (IPSEC)	46
7.3	SECURITY CONTROLS	50
7.4	NEUTRAL HOST PKI	52
7.5	OPERATOR PKI	53
7.6	GENERAL SECURITY CONSIDERATIONS	54
7.7	DoS/DDOS PROTECTION	57
7.8	GMLC LOCATION AND RADIO LOCATION LOCK	58
8	QUALITY OF SERVICE	60
9	END-TO-END MANAGEMENT	62

PARAGRAPH MARKINGS

Throughout this specification, the following paragraph markings are used:

- M** A mandatory and critical requirement that must be met by the solution. Details shall be provided stating how mandatory requirements have been met within any proposed solution.
- R** A requirement of the specification. These are to be considered mandatory to the extent that non-compliance will require the *Neutral Host* to provide to the *Operator* (or visa-versa) specific justification as to why they are not compliant to the requirement.
- I** Informative statement, providing either points of clarification or a statement relating to implementation good practice.

GLOSSARY AND ABBREVIATIONS

2FA	Two Factor Authentication
802.1p	Priority marker for (layer 2) Ethernet frames
AAA	Authentication, Authorization and Accounting (access control)
ACL	Access Control Lists
ADSL	Asymmetric Digital Subscriber Line
AES-CBC	Advanced Encryption Standard – Cipher Block Chaining
Aggregation Function	A device capable of aggregating S1, N1, N2 and N3 connections
AMF	Access and Mobility Management Function (5G core element)
AS	Autonomous System
ASCII	American Standard Code for Information Interchange (character codes)
ASIC	Application Specific Integrated Circuit
BFD	Bi-directional Forwarding Detection
BGP	Border Gateway Protocol
b -interface	The interface between the Neutral Host Domain and the Operator Domain
BTS	Base Station (e.g. picocell, eRAN cell, femtocell)
CESG	COMMUNICATIONS-ELECTRONICS SECURITY GROUP
CIR	Committed Information Rate
CM	Configuration Management
CMC	Certificate Management over CMS
CMPv2	Certificate Management Protocol version 2
CMS	Cryptographic Message Syntax
C-NAME	Canonical Name (in a DNS system)
Controller	Aggregation unit (services node) for controlling and aggregating multiple BTS
CoPP	Control Plane Policing
CP	Control Plane element (4G MME, 5G AMF)
CPE	Customer Premises Equipment (switches and routers)
CPU	Central Processing Unit
CRL	Certificate Revocation List
CU	Centralised Unit (5G)
CX600	Vendor switch
DC	Datacentre
DDoS	Distributed Denial of Service (attack)
DH	Diffie Hellman (key exchange mechanism)
DHCP	Dynamic Host Configuration Protocol
DLM	Dynamic Line Management
DNS	Domain Name Server
DNS64	DNS server which returns an synthesised IPv6 address based on hostname
DoS	Denial of Service (attack)
dot1q	VLAN tagging (layer 2)
DSCP	Differentiated Services Code Point
DU	Distributed Unit (5G)
DWDM	Dense Wavelength Division Multiplexing
eBGP	Edge BGP
EF	Expedited Forwarding
E-LINE	Ethernet (layer 2) circuit
eNodeB	4G BTS

EoFTTC	Ethernet over Fibre to the Cabinet
EPC	Enhanced Packet Core (4G Core Network)
E-RAN	Enterprise Radio Access Network
ESP	Encapsulating Security Payload
EST	Enrolment over Secure Transport
f-interface	Interface between the Retailer Domain and the Neutral Host Domain .
FM	Fault Management
FQDN	Full Qualified Domain Name
FTTC	Fibre to the Cabinet
FTTdp	Fibre to the Drop Point
FTTP	Fibre to the Premises
GGSN	Gateway GPRS service Node
gNodeB	5G BTS
GNSS	Global Navigation Satellite System (GPS or Galileo, GLONASS, BeiDou)
GPRS	General Packet Radio Service
GPS	Global Positioning System (timing source)
GRT	Global Routing Table
GTP-U	GPRS Tunnelling Protocol – User Plane
HeNB-GW	Home eNodeB Gateway
HMAC	Hash based Message Authentication Code
H-QoS	Hierarchical Quality of Service
I/C	Inter-connecting (router)
iBGP	Interior Border Gateway Protocol
IKE-SA	Internet Key Exchange – Security Association
IP	Internet Protocol (layer 3)
IPSec	IP Security protocol (encrypted packet transmission)
IPv4	Internet Protocol (Layer 3 packet switching) Version 4
ISAKMP	Internet Security Association and Key Management Protocol
IS-IS	Intermediate System to Intermediate System (routing protocol)
JOTS	Joint Operators Technical Specification (technical forum attended by UK MNOs)
LAN	Local Area Network
LNS	Local Network Service
LP	Least cost Path
MC	Mobile Core
MD5	Message Digest algorithm 5 (128 bit hash)
MEN	Metro Ethernet Network (high quality, guaranteed bandwidth)
Mgmt	Management
MGW	Media Gateway
MME	Mobility Management Entity (4G core element)
MNO	Mobile Network Operator
MOCN	Multi-Operator Core Network (a.k.a MOSS within this specification)
MODS	Multi-Operator Dedicated Spectrum (a.k.a. MORAN)
MORAN	Multiple Operator Radio Access Network (a.k.a. MODS within this specification)
MOSS	Multiple Operator Shared Spectrum (a.k.a. MOCN)
MP-iBGP	Multi-Protocol interior Border Gateway Protocol
MSC	Mobile Switching Centre (4G core element)
N1	5G interface between the UE and the AMF
N2	5G interface between the (Radio) Access Network and the AMF
N26	Interface between MME (4G) and AMF (5G)
N3	5G interface between the (Radio) Access Network and the UPF

NAPT	Network Address and Port Translation
NAT	Network Address Translation (IP layer 3)
NAT64	Network Address Translation between IPv4 and IPv6
NAT-T	NAT Traversal (UDP encapsulation)
NG	5G core N-interfaces, including (but not limited to) N1, N2 and N3
NHIB	Neutral Host In-Building
NR	5G New Radio Access Technology
NTE	Network Termination Equipment (backhaul provider)
NTP	Network Time Protocol (clock synchronisation over packet networks)
OAM	Operations and Maintenance
OCSF	Online Certificate Status Protocol
OSS	Operations Support System
OTA Sync	Over-the-Air Synchronisation (align to macrocell synchronisation)
PE	Provider Edge (router)
PIR	Peak Information Rate
PKI	Public Key Infrastructure
PLMN	Public Land Mobile Network
PM	Performance Management
pNIC	Physical Network Interface Card
PoE/+ /++	Power over Ethernet, PoE+, PoE++
ppb	Parts per billion
PTP	Precision Timing Protocol (clock synchronisation over packet networks)
QoS	Quality of Service
RAT	Radio Access Technology (i.e. 4G/5G)
RBAC	Role Based Access Control (management interface)
RIB	Routing Information Base
ROADM	Reconfigurable Optical Add Drop Multiplexor
RU	Remote Unit (5G)
S1	4G interface between eNodeB and SGW
S1-AP	S1 Application Protocol (carries user plane traffic)
S1-CP	S1 Control Protocol (carries signalling traffic)
S1-U	S1 User Plane
S5-C	Interface between Serving Gateway to PDN Gateway (control plane)
S5-U	Interface between Serving Gateway to PDN Gateway (user plane)
SA	Security Association
SCEP	Simple Certificate Enrolment Protocol
SCTP	Stream Control Transmission Protocol
SecGW	Security Gateway (terminates IPsec tunnel end points)
SGSN	Service GPRS Support Node (4G core element)
SGW	Serving Gateway (4G core element)
SHA-256	Secure Hash Algorithm (246 bit)
SIEM	Security Information and Event Management
SNMP	Simple Network Management Protocol
SR-IOV	Single Root Input/Output Virtualisation
SSH	Secure Shell protocol
SSO	Single Sign-On
SyncE	Synchronous Ethernet (transfers clock signals)
TACACS+	Terminal Access Controller Access Control System Plus
TCP	Transmission Control Protocol
Tier 1b SecGW	b -interface security gateway (within Neutral Host Domain)

Tier 1f SecGW	f -interface security gateway (within Neutral Host Domain)
Tier 2 SecGW	b -interface security gateway (within Operator Domain)
TS	Traffic Selectors (within IPSec flow)
TSCoP	Draft Telecoms Security Code of Practice
TTL	Time To Live (hop count in packet networks)
UDP	User Datagram Protocol
UE	User Equipment (typically a smartphone)
UP	User Plane element (4G: SGW, 5G: UPF)
UPF	User Plane Function (5G core element)
uRPF	Unicast Reverse Path Forwarding
UTC	Coordinated Universal Time
VLAN	Virtual Local Area Network
VNF	Virtual Network Function
VNF	Virtual Network Function
vNIC	Virtual Network Interface Card
vNIC	Virtual Network Interface Card
vRAN	Virtualised RAN
VRF	Virtual Routing Function
VSO	Vendor Specific Option (in DHCP protocol)
X2	Interface between two eNodeB (4G)
xDSL	Digital Subscriber Line (generic)
Xn	Interface between two gNodeB (5G)

1 INTRODUCTION

This JOTS specification annex sets out the architecture and central principles for the JOTS Neutral Host In-Building (NHIB) concept. For brevity, these will be referred to as the JOTS NHIB concept or JOTS NHIB specification from this point onwards.

It introduces the **Retailer Domain**, **Neutral Host Domain** and **Operator Domain** and explains how these sit relative to each other within the overall architecture. Responsibilities within each domain are defined along with the connectivity between the domains.

Key to the implementation of the JOTS NHIB concept are the routing principles and security models which underpin the method by which multiple *Operators* can share common aggregation infrastructure. The specific design concepts, important mandatory security requirements and control measures are set out in full in this specification.

The JOTS NHIB specification is expected to be an evolving specification which will be updated as and when required (by the JOTS forum) to maintain alignment and relevance with new technologies and developing vendor capabilities.

This version of the specification no longer differentiates between non-virtualised and virtualised implementations. Both approaches are supported for rollout, although some capacity or throughput limitations typical in some virtualised components must be properly considered where a virtualised approach is used.

IP addressing towards the *Operator's* core networks will be based on IPv4 for initial deployments, however a 3-step evolution path towards IPv6 support is detailed within this specification. It should be assumed that each *Operator* will evolve their IP addressing independently, meaning the *Neutral Host* will be required to support IPv4 and the three stages of IPv6 evolution concurrently.

The JOTS NHIB specification doesn't prescribe or prefer non-virtualised implementations over virtualised ones. Nor does it prefer any vendor solution over any other. The aim of the specification is to define a set of requirements, which, if met, will enable an NHIB platform to be deployed in a *Neutral Host* datacentre and connected to one or more mobile core networks concurrently.

The JOTS NHIB specification supports various BTS deployment models across various types of venues. Again, the specification does not prefer any one deployment model over any other.

2 DOMAINS

- 1.1 The Neutral Host In-Building (NHIB) deployment is separated into three domains: the **Retailer Domain**¹, the **Neutral Host Domain** and the **Operator Domain**, where the key areas of responsibility are shown in *Figure 2-1*:

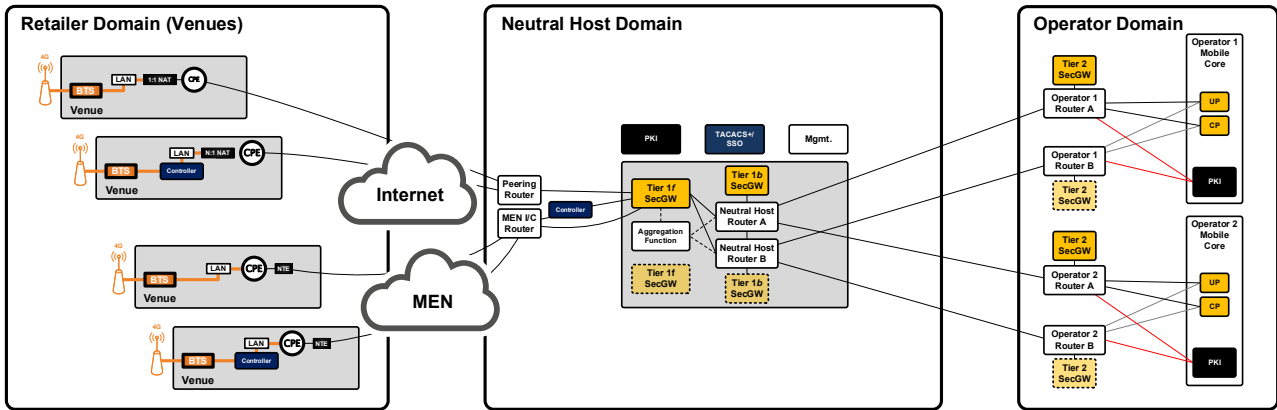


Figure 2-1 - Domain Overview.

A number of possible radio architectures are achievable depending on the needs of the selected radio solution and on the aggregation requirements of the **Operator Domain**. Example radio architectures are shown in *Figure 2-2*, *Figure 2-3*, *Figure 2-4*, *Figure 2-5* and *Figure 2-6*. For illustrative purposes, the wavy line depicts the traffic path adopted for each radio architecture option.

A *Neutral Host* would select and operate one or perhaps multiple radio architectures, however it is not necessary for a *Neutral Host* to implement *all* types of radio architecture.

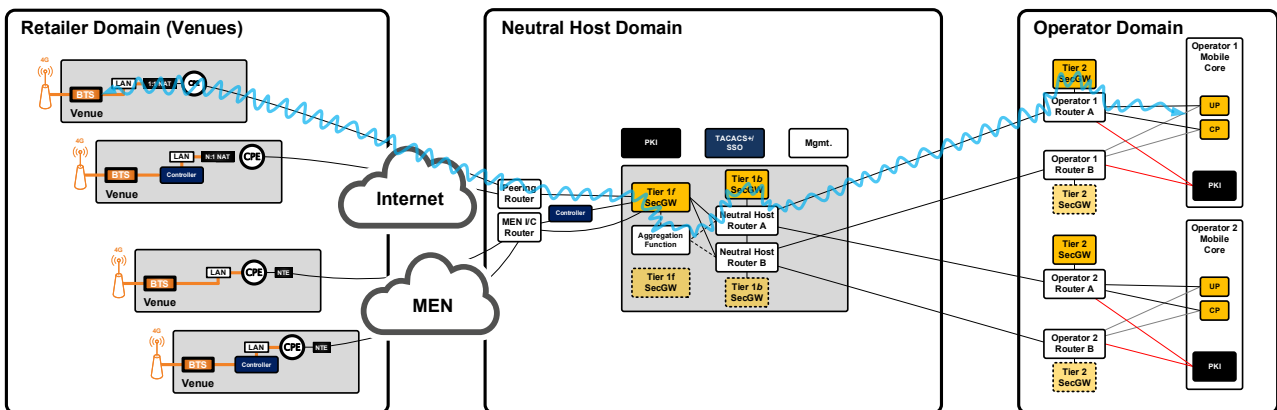


Figure 2-2. Internet f-interface with traffic aggregation in Neutral Host (Aggregation Function) Domain only.

¹ For the avoidance of doubt, a *Retailer* within the **Retailer Domain**, in this context, is not a ‘shop’, but an entity whose commercial model is built around providing in-building coverage solutions or deploying small cells.

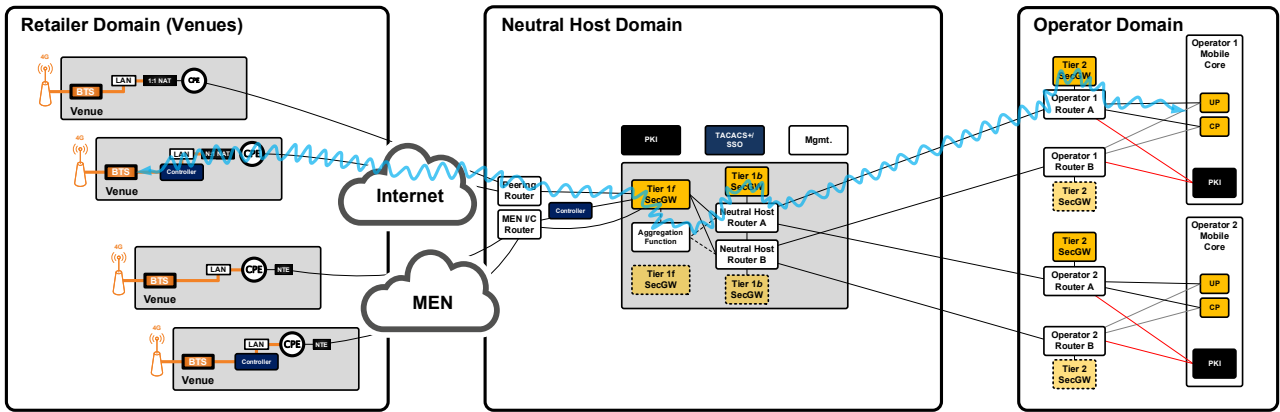


Figure 2-3. Internet f-interface with traffic aggregation in both Retailer (Controller) and Neutral Host (Aggregation Function) Domains.

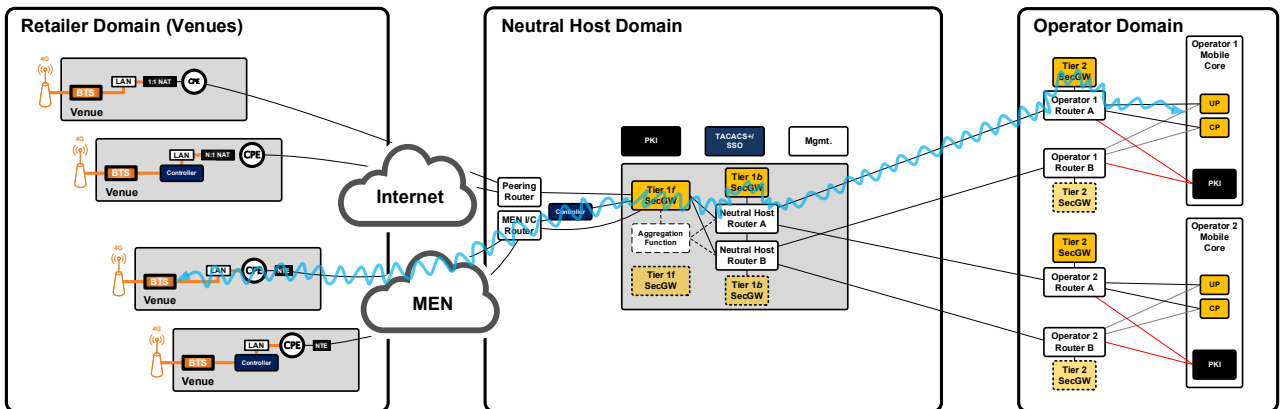


Figure 2-4. MEN f-interface with traffic aggregation in Neutral Host (Controller) Domain only.

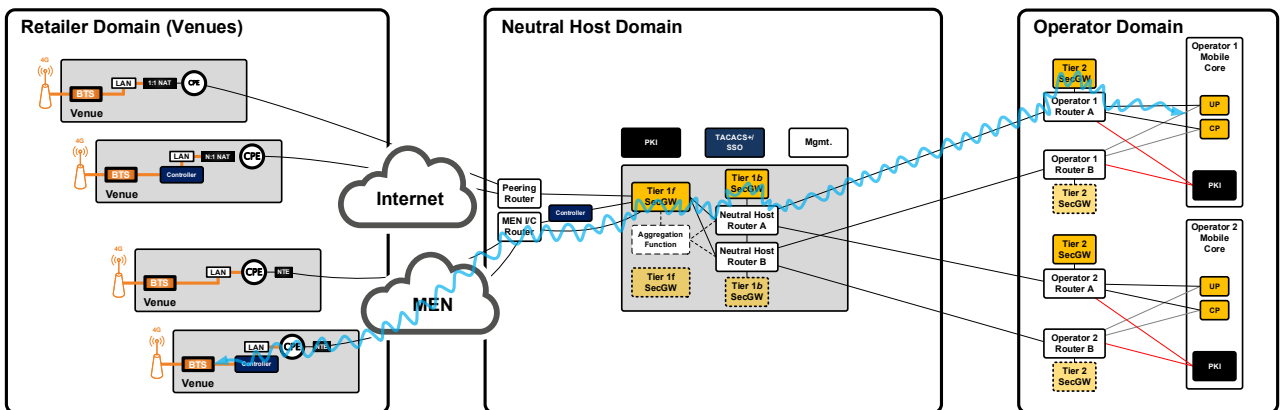


Figure 2-5. MEN f-interface with traffic aggregation in Retailer (Controller) Domain only.

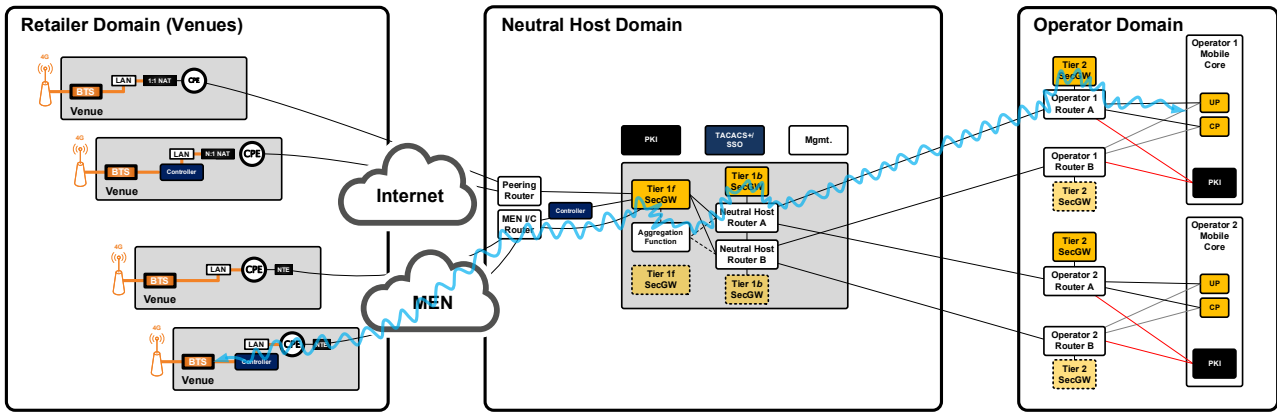


Figure 2-6. MEN *f*-interface with traffic aggregation in both Retailer (Controller) and Neutral Host (Aggregation Function) Domains.

2. I The following high-level statements can be made to describe the **Retailer Domain** and the responsibilities of the *Retailer* (who operates within this domain):

- That procurement and provision of site equipment (i.e. BTS, Controllers, supporting switching infrastructure, etc.) is the responsibility of the *Retailer*;
- That the *Retailer* is responsible for the relationship with the end customer (i.e. venues²) and their connectivity (referred to as *f*-interface) to the *Neutral Host* in their datacentre environment;
- That the *Retailer* is responsible for the configuration of the site equipment (including associated customer site related security measures);
- That site assessment, *f*-interface pre-qualification, post-qualification, monitoring, capacity management and upgrade assessment, reporting, governance, change management, operations, trouble ticketing etc. all need agreement and commitment from the *Retailer* for the components for which they are responsible.

3. M The *Neutral Host* must maintain documentation relating to any *Operator* specific requirements, in respect of design decisions or architectural principles which must be adhered to, and must provide evidence to the *Operator*, when requested to do so, that these *Operator* specific requirements have been met by the *Neutral Host* and have been passed onto the *Retailer*³ for implementation.

4. I The following high-level statements can be made to describe the **Neutral Host Domain** and the responsibilities of the *Neutral Host* (who operates within this domain):

- That the Tier-1*f* SecGW components (terminating BTS or Controller initiated IPsec tunnels) are deployed in the chosen datacentre of the *Neutral Host*;
- That appropriate connectivity from the datacentre of the **Neutral Host Domain** to the **Operator Domain** is in place, such that connectivity can be achieved to the Tier-2 SecGW components (terminating *Operator* specific tunnels from the Tier-1*b* SecGW);

² The term ‘venue’ is used throughout this annex to refer to the physical location where the small cell BTS is deployed.

³ It is perfectly feasible for the *Retailer* and *Neutral Host* to be business functions provided a common commercial entity.

- That the *Neutral Host* will be required to provide an Aggregation Function of BTS endpoints in order to minimise the scale of onward presentation to the *Operator* mobile core elements (e.g. MME or AMF)⁴;
- That the *Neutral Host* will be required⁵ by some or all *Operators* to provide a MOSS capability in order to support multi-PLMN working on a single carrier. Thus, the *Neutral Host* will need an ability to steer traffic for unique PLMNs to specific core network 'pools' within a single **Operator Domain**. This may be realised as part of the chosen radio solution where such capability exists, or by means of a centralised aggregation/steering function.
- That RBAC controlled multi-tenant management, reporting and CM/FM/PM⁶ capability will be required at the platform within the **Neutral Host Domain**, and that the platform will provide visibility (read access) on a per *Operator* basis. Where an *Operator* requires direct (write) access to the management platform, the specific access privileges and configurable parameters will be mutually agreed between the *Operator* and the *Neutral Host*;
- In-band or out-of-band management connectivity from the *Neutral Host* towards each participating *Operator* must be in place;
- That appropriate PKI is in place to support authentication of BTS tunnel instantiation;
- That appropriate ISO27001 and draft TSCoP⁷ governance measures are met by the *Neutral Host* in order to not compromise these requirements within the overall deployment of each *Operator*.

5.1 The following high-level statements can be made to describe the **Operator Domain** and the responsibilities of the *Operator*:

- That an appropriate Tier-2 SecGW (or equivalent IPsec capable device/function⁸) is deployed in order to terminate the tunnels from the *Neutral Host* Tier-1**b** SecGW;
- That appropriate PKI is in place to support authentication of Tier-1**b** to Tier-2 SecGW tunnel instantiation;
- That onward connectivity from the Tier-2 SecGW towards the relevant mobile core elements is provided.

6.1 In terms of connectivity responsibility, the following can be stated:

- That within a domain, the domain owner is responsible for detailed design and realisation of the connectivity;
- That at the interfaces between domains, a level of joint responsibility is present but that for key items an owner must be defined and agreed within an interface contract.

⁴ This Aggregation Function may exist in a number of different forms and can occur before the Tier-1**f** SecGW, between the Tier-1**f** and Tier-2 SecGW functions, or a combination of both, depending on the deployed technology and the level of aggregation required by the **Operator Domain**.

⁵ It should be assumed by *Neutral Hosts* that dual-PLMN-per-*Operator* support will be required in order to achieve full coverage of *Operators*, given that at least one *Operator* will provide dual-PLMN connectivity using distinct PLMNs and distinct mobile core pools. More generally, future use of MOSS across the *Operators* for in-building solutions can be expected to move this to a general Requirement (R) for all *Operators* in a future release of this specification.

⁶ Configuration Management (CM), Fault Management (FM), Performance Management (PM).

⁷ Draft Telecommunications Security Code of Practice and Telecommunications (Security) Act 2021 (legislation.gov.uk).

⁸ The term 'SecGW' is used for the Tier-1**b** and Tier-2 perimeter functions within this annex, given the expectation that this function will typically be served by a dedicated context on an existing SecGW device within the **Operator Domain**. However, these functions do not need to be a SecGW in the formal sense, they simply need to be devices capable of adhering to the IPsec requirements outlined within this annex.

3 CONNECTIVITY

3.1 Venue Connectivity

7. I Venue connectivity can take several detailed forms and it is not the intended scope of this annex to cover all feasible variants, but instead to highlight key requirements and typical deployment components.
8. I At the venue it is typically expected that the BTS devices will be connected via an existing switched environment to a local CPE, but Layer 2 (VLAN extension) services from the venue to the **Neutral Host Domain** can also be supported.
9. I To simplify deployment, it is recommended that at larger venues an appropriate PoE switch capability (i.e. correct variant of PoE/+/>++ to meet required power levels) be deployed for provision of initial hop connectivity from the BTS devices. Where replacement/augmentation of the existing switch infrastructure with PoE capable switches is impractical, PoE injectors are admissible, but it must be noted that these cannot be practically managed or monitored (if at all).
10. R Since it is anticipated that switch infrastructure at the venue *might*⁹ be shared with other switched traffic, it is required that all switch ports use dot1q encapsulation (including the router port facing the LAN), such that prioritisation based on 802.1p marking can occur where required and allow appropriate traffic handling should congestion occur on the LAN.
11. R Where a CPE and Layer 3 domain is deployed at the venue, every venue will require a local IP network to handle traffic between the BTS devices and the CPE, with the size of the network clearly dependant on the planned number of BTS nodes. Where a Layer 2 extension service is used between the venue and the **Neutral Host Domain**, similarly an appropriately 'venue-aligned' IP network allocation will need to be made per venue. It is recommended that venue subnet size allocation is made in such a way to allow for some growth without having to re-address the network. Note that this addressing is used for transport (i.e. for creating the tunnel-outer) and is therefore typically either re-usable (across different venues) private addressing undergoing 1:1 or N:1 (overloaded) NAT to a public IP, or dedicated private addressing to provide connection via Metro Ethernet (i.e. VLAN extension) type services between the venue and the **Neutral Host Domain**.
12. I Where a CPE is deployed, it is expected that the CPE will host a DHCP server function in the **Retailer Domain** (i.e. at the venue), in order to provide pool-based outer-address allocation for the BTS nodes, with VSO (Vendor Specific Option, DHCP Option 43) or other custom options where necessary to provide connectivity information for the BTS beyond its own venue pool address allocation (e.g. to a Controller node). Where a Layer-2 extension service to the **Neutral Host Domain** is used, it is expected that the DHCP function (if used) will be provided within the **Neutral Host Domain**.

⁹ Typically dedicated switched environments are deployed in larger venue deployments. These can be considered to reduce the reliance on 802.1p capability in the LAN (at least if sufficiently dimensioned to avoid congestion), but it remains best practice to mark appropriately in any case.

13. R For cases where the *f*-interface IPsec tunnel does not connect to a tunnel endpoint resolved via DHCP VSO, resolution of tunnel endpoint at the Tier 1*f* SecGW in the **Neutral Host Domain** should be achieved wherever possible by FQDN (Fully Qualified Domain Name) resolution.
14. M The venue must have DNS resolution capability available to the small cell service, either by means of local DNS¹⁰ (with appropriate A-record and C-record population), **Neutral Host Domain** hosted DNS, or appropriately secured connectivity to public DNS where the appropriate FQDN must be present.
15. R It is strongly recommended that a specific VRF instance is created at the CPE for the small cells service, hosting the venue addressing pool and DHCP server where required.
16. I From a high-level point of view a BTS will be setup as follows:
- a. The BTS will be configured with the required VLAN encapsulation for connectivity through the switched network towards the small cells VRF at the CPE;
 - b. At this point the BTS has access to the local network, it will then discover basic connectivity information:
 - IP address, network and default gateway via standard DHCP request;
 - Controller IP address where appropriate - typically obtained via DHCP option 43 (or provided by appropriate DHCP options for the deployed equipment type);
 - Alternatively, tunnel endpoint addresses should be determined via FQDN.
 - c. Once the small cell has an IP address and termination information it will establish a connection either to the Controller or the Tier 1*f* Security Gateway:
 - In order to establish the IPsec tunnel (i.e. the IKE-SA), it will be necessary for certificated authentication to take place (see PKI section for more details);
 - Once the tunnel is established, OAM connectivity is expected to be achieved using the same tunnel-inner address used for S1-AP/S1-U/NG connection¹¹. It should be noted that where specific endpoints are deployed for different MNOs, management segregation is required.

¹⁰ Clearly any local DNS deployment must have an appropriate in-life process to ensure its content is up to date.

¹¹ Where the deployed equipment requires a separate tunnel for OAM connectivity, this can be considered.

3.2 Venue Types

17. I The two most typical venue connectivity models, Metro Ethernet Network (MEN) and Internet, along with NAT variants in the Internet connectivity model are considered. Each is discussed with differing Controller and Aggregation Function presence and location.

18. I Venue type A is illustrated in *Figure 3-1*:




Figure 3-1 - Venue Type A.

- BTS connected via LAN switch to CPE VRF via 1:1 static symmetric NAT function;
- CPE to provide DHCP server within the VRF, with appropriate addressing pool for tunnel-outer address allocation and Vendor Specific Option (VSO) configuration as appropriate;
- 802.1p support in the LAN¹²;
- DSCP and 802.1p marked packets at the BTS;
- Internet bearer (which may be either Ethernet or xDSL based Internet access). (It should be noted that for a centralised Controller type deployment which typically has a requirement for low jitter and latency *between* the BTS and Controller function, it should be considered that ADSL variants are not capable of supporting the requirements and therefore an FTTC or FTTP service will be required if xDSL is used);
- Centralised Controller function (e.g. E-RAN Service Node);
- Additional Aggregation Function where necessary to achieve the level of aggregation required by the **Operator Domain**.

¹² 802.1p support will not be strictly required in the cases where dedicated LAN infrastructure is provided for the BTS connectivity and where sufficient capacity exists in that dedicated LAN environment such that congestion does not occur. 802.1p support is included since it remains best practice in order to maintain prioritisation of traffic classes if and when congestion does occur.

19.1 Venue type B is illustrated in *Figure 3-2*:




Figure 3-2 - Venue Type B.

- BTS connected via LAN switch to CPE VRF via N:1 overloaded NAPT function;
- CPE to provide DHCP server within the VRF, with appropriate addressing pool for tunnel-outer address allocation and Vendor Specific Option (VSO) configuration as appropriate;
- 802.1p support in the LAN;
- DSCP and 802.1p marked packets at the BTS;
- Internet bearer (which may be either Ethernet or xDSL based Internet access). (Noting that due to the co-located Controller and the less stringent constraints on connectivity between the Controller and Tier-1f SecGW, ADSL services could *potentially* be considered acceptable here as well as FTTC and FTTP);
- Co-located E-RAN Controller (i.e. Service Node);
- Additional Aggregation Function where necessary to achieve the level of aggregation required by the **Operator Domain**.

20.1 Venue type C is illustrated in *Figure 3-3*.




Figure 3-3 - Venue Type C.

- This venue type may be deployed as a Layer 2 VLAN extension between BTS and **Neutral Host Domain** and as such the CPE component can be considered optional;
- In the case where CPE is provided (with a local L3 domain):
 - BTS connected via LAN switch to CPE VRF without NAT;
 - CPE to provide DHCP server within the VRF, with appropriate addressing pool for tunnel-outer address allocation and VSO configuration as appropriate;
- In all cases:
 - 802.1p support in the LAN;
 - DSCP and 802.1p marked packets at the BTS;
 - Ethernet bearer (noting that this will typically be via a MEN E-LINE type service);
 - No Controller function;
 - Optional additional Aggregation Function.
 - Additional Aggregation Function where necessary to achieve the level of aggregation required by the **Operator Domain**.

21.1 Venue type D is illustrated in *Figure 3-4*:




Figure 3-4 - Venue Type D.

- This venue type may be deployed as a Layer 2 VLAN extension between BTS and **Neutral Host Domain** and as such the CPE component can be considered optional;
- In the case where CPE is provided (with a local L3 domain):
 - BTS connected via LAN switch to CPE VRF without NAT;
 - CPE to provide DHCP server within the VRF, with appropriate addressing pool for tunnel-outer address allocation and VSO configuration as appropriate;

- In all cases:
 - 802.1p support in the LAN;
 - DSCP and 802.1p marked packets at the BTS;
 - Ethernet bearer (noting that this will typically be via a MEN E-LINE type service);
 - Co-located Controller function;
 - Optional additional Aggregation Function.
 - Additional Aggregation Function where necessary to achieve the level of aggregation required by the **Operator Domain**.

3.3 *f*-interface Connectivity

22. R *f*-interface connectivity responsibility lies with the *Retailer* within the **Retailer Domain**.
23. R *f*-interface connectivity can be provisioned by means of one or more of the following:
- Private Ethernet services (i.e. VLAN extension or E-LINE type services);
 - Private xDSL services (if the **Neutral Host Domain** has the appropriate equipment e.g. LNS and interconnects with an access provider to provide this type of service to the **Retailer Domain**);
 - Public Internet Ethernet services (i.e. a symmetric Ethernet Internet access service);
 - Public Internet xDSL or fibre broadband services.
24. M The **Neutral Host/Retailer Domains** will be responsible for selecting an appropriate transport technology to meet the bit rate, latency, jitter and packet loss performance requirements of the *f*-interface component of the radio solution.
25. I Where xDSL or fibre broadband services are used, the effect of contention and lack of QoS features (noting that some limited elevated product features might exist depending on the provider and product selection) must be properly considered when qualifying a site.
26. I For an Internet-based *f*-interface, it can be expected that one of the following approaches is used:
- Transport using shared capacity on an existing Internet service already provided at the venue;
 - Capacity on an additional Internet service at the venue procured specifically for use by the BTS-footprint.
27. R Where an existing Internet provision is shared, qualification and ongoing monitoring in the **Retailer Domain** must ensure that capacity and performance of the service remains sufficient to provide the expected customer experience (noting that qualification performed solely at specific times of the day will not be representative of the fluctuating loads of a shared Internet service). Evidence supporting the qualification and ongoing monitoring should be provided by the *Retailer* to the *Neutral Host* and be made available to the *Operator*.

28. R Where xDSL services are deployed, qualification and ongoing monitoring is the responsibility of the *Retailer* in the **Retailer Domain**. Initial qualification of the *f*-interface (post deployment) against the requirements of the service must consider the initial management cycle of Dynamic Line Management (DLM), given that this will typically vary the line configuration and consequently the performance characteristics quite aggressively over an initial 10-day period and may remain active in life to cater for seasonal variation in copper plant performance. Evidence supporting the qualification and ongoing monitoring should be provided by the *Retailer* to the *Neutral Host* and be made available to the *Operator*.
29. R Should in-life monitoring demonstrate sub-standard customer experience or capacity issues, the following lifecycle of upgrades shall be actioned within the **Retailer Domain** for affected venues where the condition is not shown to be the result of other (resolvable) faults:
- Where shared Internet *f*-interface services are used:
 - Resolution step 1 - Addition of BTS-footprint dedicated Internet service;
 - Where dedicated Internet *f*-interface services are used:
 - Resolution step 1 - Upgrade of BTS-footprint dedicated Internet service;
 - Resolution step 2 - Upgrade to committed bandwidth private Ethernet service.
 - Where private Ethernet *f*-interface services are used:
 - Resolution step 1 - Upgrade to higher capacity private Ethernet service.
 - Where public or private xDSL *f*-interface services are used:
 - Resolution step 1 – Upgrade to shared or dedicated Internet service or private Ethernet service.
30. R Wherever public Internet services are used (and potentially where private Ethernet services are used), IPsec tunnel establishment for the *f*-interface will need to take place via a private to public Network Address Translation (NAT) function. This NAT capability can take one of the following forms:
- NAT (Static 1:1 symmetric NAT);
 - NAPT (N:1 overloaded NAT).
31. I In order to transport the Encapsulating Security Payload (ESP) traffic, which cannot by itself traverse NAPT (N:1 overloaded NAT) due to a lack of Layer 4 information, this will require NAT-Traversal by encapsulation in UDP Port 4500. Whilst this mechanism will be automatic, appropriate consideration must be made to allow for this in any firewall rules within the **Retailer Domain** or **Neutral Host Domain**¹³.

¹³ It should be noted that synchronisation related traffic (e.g. NTP) may be outside of the IPsec tunnel, depending on the deployed technology. If this is the case, appropriate firewall rules will clearly need to be made to cater for this.

32. M Synchronisation services must be provided that are appropriate for the deployed radio solution requirements, with these typically expected to use NTP, PTPv2¹⁴, SyncE, GNSS and OTA or in some cases a combination of these technical approaches.
33. M It is mandated that timing services are provided from within the **Neutral Host Domain**.
34. M Components in the **Retailer Domain**, **Neutral Host Domain** and **Operator Domains** must all derive an accurate time-of-day level clock in order that accurate (and therefore comparable) timestamping of events (configuration changes, logs etc.) and validity of certificates is accommodated.
35. R Where a shared Internet provision is used at the venue the BTS footprint must be able to make use of NAPT (N:1 overloaded NAT).
36. I Where a dedicated Internet provision is used the BTS-footprint can make use of either NAT or NAPT, dependant on how many distinct addresses are required at the venue, noting that NAT is simpler since ESP NAT traversal is not an issue.

3.4 **b-interface Connectivity**

37. R **b-interface** connectivity responsibility lies primarily with the *Neutral Host* within the **Neutral Host Domain**.
38. R **b-interface** connectivity can be provisioned by means of one or more of the following:
- A Dense Wavelength Division Multiplexing (DWDM) connection ('DWDM transmission');
 - A Metro Optical Network connection ('Metro transmission');
 - A dedicated Point-to-Point Layer 2 Service connection ('E-LINE transmission');
 - An Internet connection via Private Peering, Public Peering or Transit Peering ('Internet transmission').

¹⁴ To support TDD on the radio interface (which includes 5G) the accuracy of PTPv2 G.8275.1 time synchronisation, relative to a standard time reference (e.g. GNSS) measured at the input to the base station, must be better than $\pm 1.1 \mu\text{s}$ in order to achieve a cell synchronisation better than $\pm 1.5 \mu\text{s}$ relative to the standard time reference.

39. I The **b**-interface connectivity options are illustrated in *Figure 3-5*:

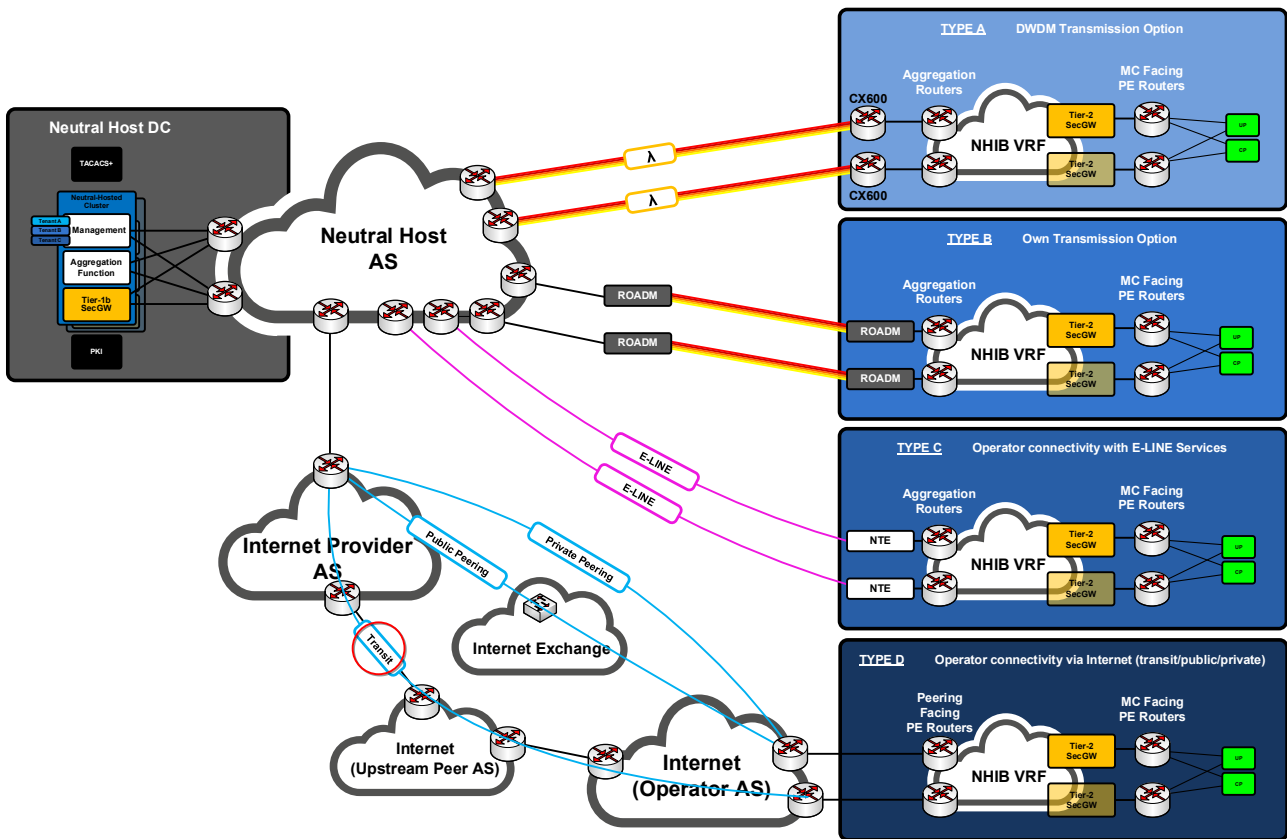


Figure 3-5 - Neutral Host to Operator Connectivity Options.

40. I The **b**-interface connectivity option implemented between the **Neutral Host Domain** and each **Operator Domain** does not necessarily have to be of the same type. The connectivity option will be selected by mutual agreement between the *Neutral Host* and each *Operator* separately. Furthermore, a single *Neutral Host* might mutually agree to implement multiple connectivity option types towards a single *Operator*.
41. R The requirement for the **b**-interface connectivity to be a QoS-enabled private connectivity service will be the decision of each *Operator* against their QoS and SLA requirements and is expected to be both the typically deployed option and the required option upon reaching certain points of scale. Internet backhaul can be supported where required for initial deployments and for secondary paths until such point as the *Operator* subject to that connectivity deems it no longer appropriate for the required scale and level of service availability/quality.

4 ROUTING

4.1 Routing Options

42. I Traffic is separated at all points in the solution where it is practical to do so, but it must be noted that for a non-virtualised implementation there are some points where the S1/NG traffic for all *Operators* converges.
43. I For the purpose of traffic segregation (and therefore security), routing from the Tier-1f SecGW or Aggregation Function will be presented to *Operator* specific VRFs in the **Neutral Host Domain**, from where routing towards the *Operator* will be established.
44. R Two routing options are permitted for the *b*-interface connection between the **Neutral Host Domain** and the **Operator Domain**:
- **Option 1**: assumes the Tier-1f SecGW function or the Aggregation Function in the **Neutral Host Domain** routes traffic directly towards *Operator* specific VRFs;
 - **Option 2**: assumes the Tier-1f SecGW function or the Aggregation Function in the **Neutral Host Domain** is *only* able to route traffic towards *Operator* specific VRFs via a common Landing VRF.
45. I It is expected that **Option 1** will become *mandated* in a future version of the specification. However, it is accepted that the capability to achieve this depends on the supported functionality of the selected *Neutral Host* equipment and cannot therefore be mandated initially. Where interface/sub-interface separation from the Tier-1f SecGW or Aggregation Function is not possible, **Option 2** provides a workable solution *for an interim period*.

46. I An overview showing the **Option 1** approach is illustrated in *Figure 4-1*:

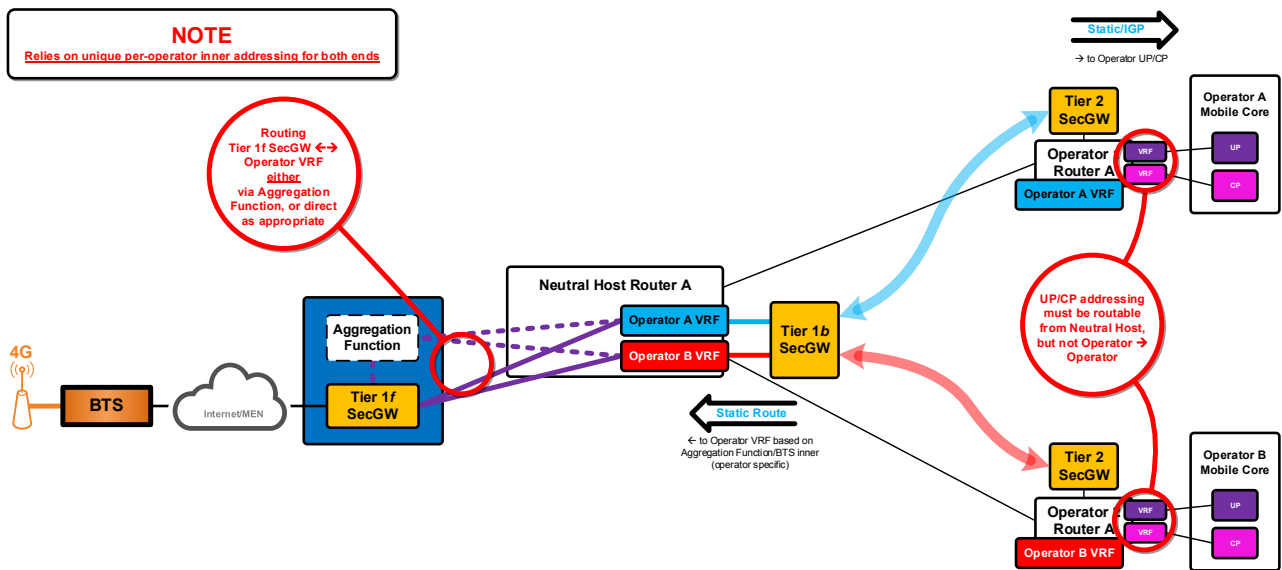


Figure 4-1 – Option 1 Routing Approach.

47. I An overview showing the **Option 2** approach is illustrated in *Figure 4-2*:

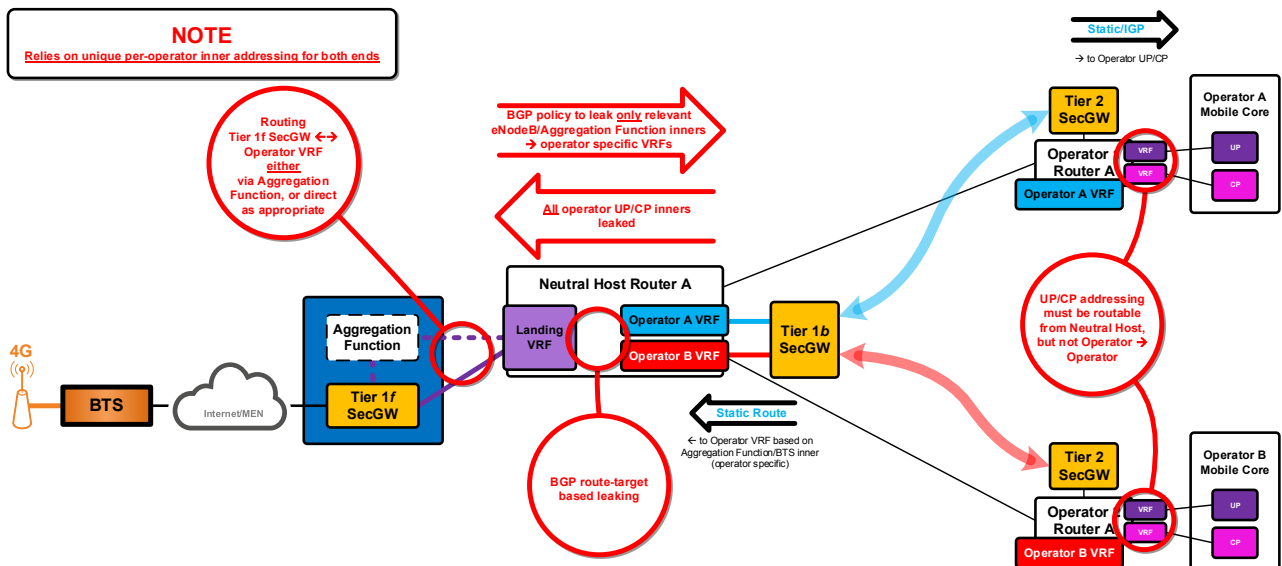


Figure 4-2 – Option 2 Routing Approach.

48. R **b**-interface traffic shall be delivered to the *Operator* specific VRFs from the Tier-1**f** SecGW or Aggregation Function by the following means:

- **Option 1:** Multiple physical (or multiple dot1q encapsulated logical on a shared physical) connections, one per *Operator*, from the Tier-1**f** SecGW or Aggregation Function towards

Operator specific VRFs at the *Neutral Host* router, with unique IP per-*Operator* mobile core destinations used to route via the correct exit interface;

- **Option 2:** Single physical (or dot1q encapsulated logical) connection for all *Operators* from the Tier-1f SecGW or Aggregation Function towards *Operator* specific VRFs at the *Neutral Host* router, with BGP route-target based route leaking within the Landing VRF (ideally towards a distinct PE node) into unique per-*Operator* VRFs based on applied (and controlled) BGP policy.

49. M For **Option 2** it should be noted that route leaking between VRFs and the Global Route Table (GRT) will not be used (nor accepted), due to the potential security risks of this approach, particularly as a result of in-life misconfiguration.
50. M For **Option 2** it is essential that appropriate BGP policies are in place to constrain leaked prefixes to those required only to achieve transport between the **Neutral Host Domain** and **Operator Domain** and to prevent reachability between *Operator* networks. It is essential that default routes are never sent in either direction and that all policies must specifically ensure this¹⁵.
51. I For evolving RAN capabilities, such as vRAN, it can be considered that rather than sharing components in the **Neutral Host Domain**, these could become per-*Operator* dedicated instances (typically virtualised).
52. I Per-*Operator* dedicated instances (typically virtualised) reduce or remove the convergence of traffic for multiple *Operators*¹⁶, and simplifies IP addressing between *Operators* (as these can be contained within per-*Operator* VRFs and SecGW TS instances, with addressing overlap) but does not in fact change the routing, resilience or security mechanisms covered within this specification.
53. I Per-*Operator* instances can be shared or discrete, as shown in *Figure 4-3*. It is recommended that consideration of the requirements of this model is made during equipment selection to cater for the radio evolution path potentially dictating discrete instances per *Operator* in the future. Where they are discrete, routing and IPsec policy (e.g. Traffic Selectors) become compartmentalised for each *Operator* and therefore can be overlapped. Security policy could conceptually be simplified in this model as the end-to-end logical separation of routing further reduces the possibility of inter-*Operator* routing. But in any case, the security controls outlined in this specification must still be applied for the additional protection against misconfiguration.

¹⁵ As detailed later in the annex, routing policy must in fact constrain advertisements to *only* those specifically required prefixes in either direction.

¹⁶ At least in the MORAN/MODS model.

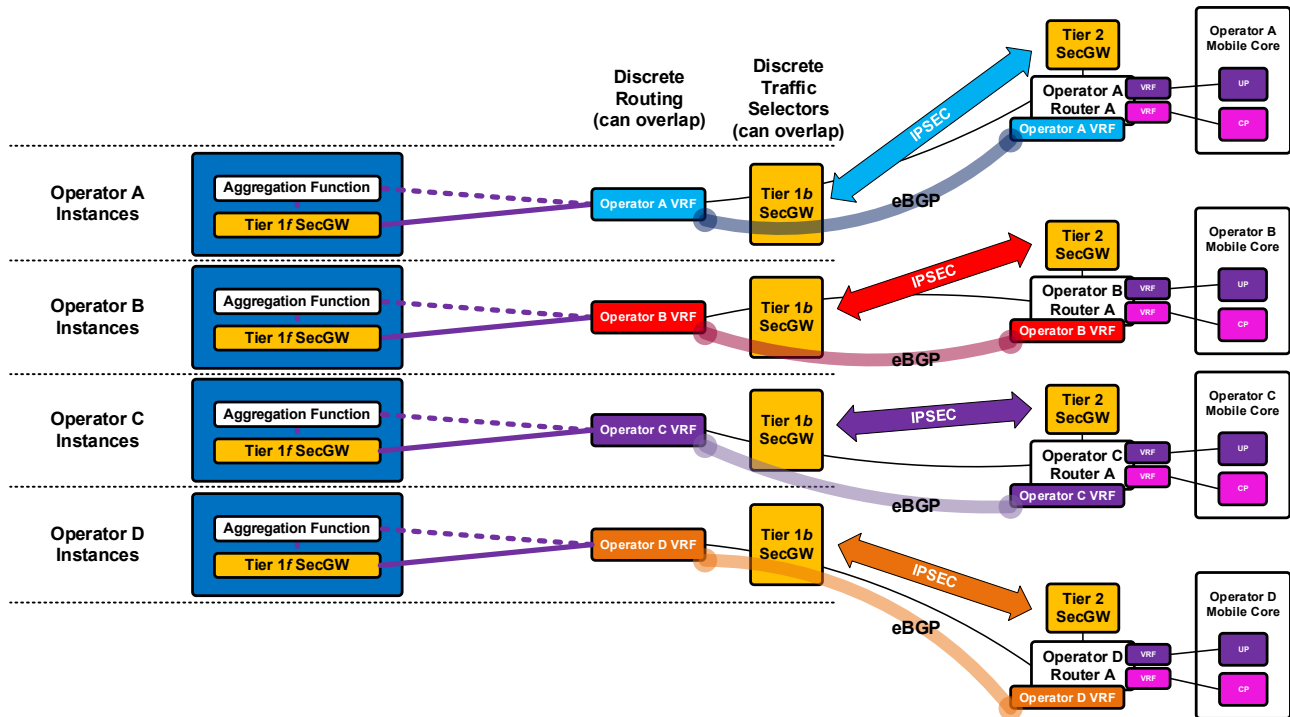


Figure 4-3 – Per-Operator Instances (typically virtualised).

54. I IPsec tunnels between *Neutral Host* Tier-1**b** SecGW and *Operator* Tier-2 SecGW are built in the same manner, along with eBGP established through each tunnel to the respective *Operator*, albeit now from a discrete-per-*Operator* Tier-1**b** SecGW rather than a shared device. Similarly, if the Tier-1**f** SecGW and Aggregation Function (if present) are discrete-per-*Operator* instances, routing between these devices and the Tier-1**b** SecGW can also be segregated at VRF level.
55. M For *b*-interface connectivity, regardless of whether the connectivity is achieved in the public (i.e. Internet) or private network domains, the use of IPsec is mandated.
56. I The use of IPsec on the *b*-interface connection provides:
- A standardised certificated authentication model to support establishment of connectivity (and lifecycle);
 - An encryption model, where necessary, without presenting different design and configuration requirements dependent upon the type of *b*-interface connectivity selected (effectively providing a *b*-interface agnostic approach¹⁷).
57. M IPsec connectivity must be established between the **Neutral Host Domain** and each **Operator Domain** independently.

¹⁷ The intention here is also that at such point as a *Neutral Host* ‘re-grades’ the *b*-interface connectivity from an Internet to a private bearer (or vice-versa), the routing and security model does not change.

- 58. M From a routing perspective all traffic between the **Neutral Host Domain** and the **Operator Domain** will transit an established IPsec tunnel and that this IPsec connectivity must support transport between several IP subnet pairs.

- 59. I Routing onwards from the Tier-2 SecGW component towards the *Operator* mobile core devices is not the subject of this annex. It is expected that this is typically achieved by means of multiple sub-interfaces from the SecGW bound into the relevant VRFs at the *Operator* PE nodes, with exit sub-interfaces selected purely based on destination IP, or via a Landing VRF, which itself has the necessary onward routing paths. However, regardless of the method, this connectivity is a per-*Operator* decision.

4.2 eBGP Peering

- 60. I A summary of the eBGP peering (in a non-resilient model¹⁸) is illustrated in *Figure 4-4*:

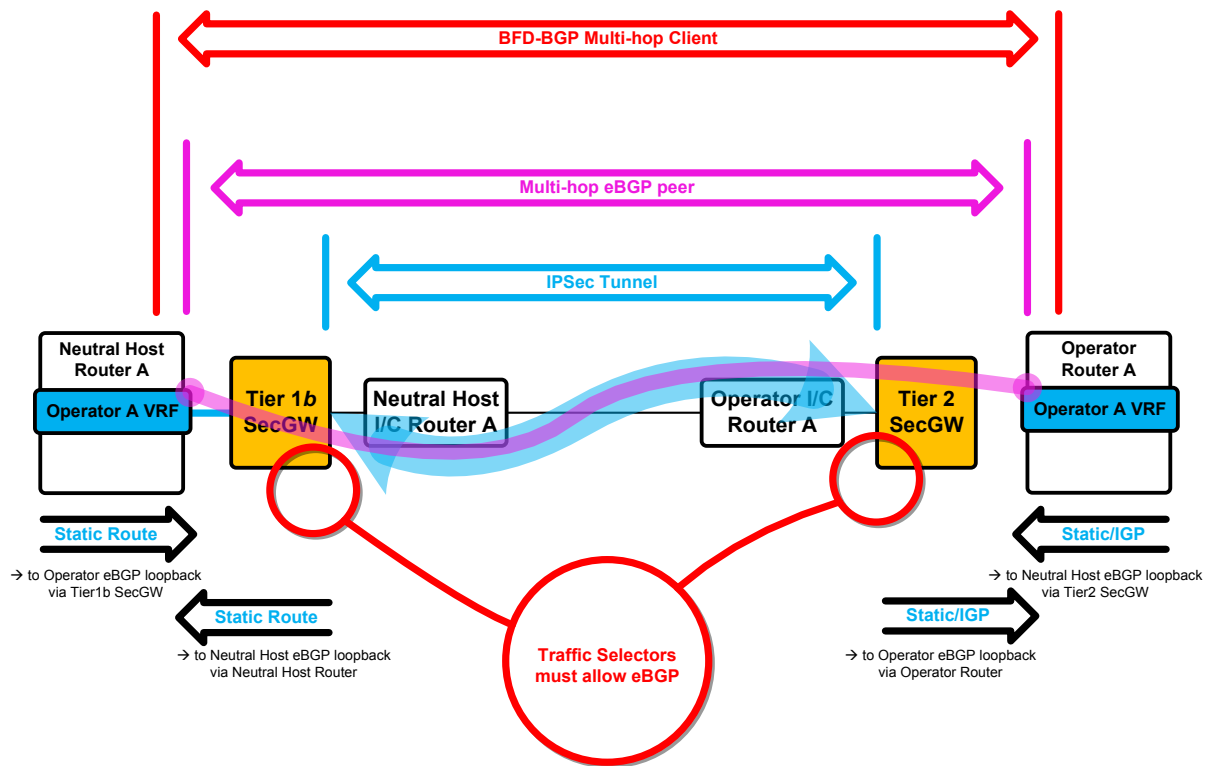


Figure 4-4 – b-interface eBGP Peering Approach.

- 61. I Given the base connectivity model shown in *Figure 4-1* (Option 1) and *Figure 4-2* (Option 2), it is clear that establishment of eBGP for routing purposes is not a simple direct peering connection, as would typically be the case for eBGP. Instead, it will be necessary for the eBGP peering to be

¹⁸ The non-resilient model is shown for the purposes of explaining the eBGP peering principles.

established between the *Neutral Host* router and *Operator* router (with this established from each appropriate *Operator* specific VRF), *via* the Tier-1**b** and Tier-2 Security Gateways.

- 62. I For both eBGP and S1-AP/S1-U/NG transport, given that the routing path is indirect (i.e. the BGP next-hop must be recursed to be reachable) and given that the intermediate Tier-1**b** SecGW and Tier-2 SecGW are not taking part in any dynamic routing protocol exchanges), it is necessary for appropriate routes to exist at those nodes as well as the eBGP learned routes at the PE nodes.
- 63. I For the **Neutral Host Domain** side, it is expected that the Tier-1**b** SecGW is both physically and logically close to the *Neutral Host* router and as such it is expected that static routing will be used for this component.
- 64. I In the **Operator Domain**, it is expected that either static or IGP distributed routes will be used to complete reachability towards the **Neutral Host Domain**.
- 65. I Reachability to the eBGP peering loopback must be achieved indirectly. It is expected that the PE router hosting the eBGP session towards the other party (in both the **Neutral Host Domain** and **Operator Domain**) will be 'close' (i.e. either the same device or directly connected) to the Tier-1**b** or Tier-2 SecGW function and that static routing will therefore be appropriate to provide an initial hop via the SecGW towards the far-end eBGP peer.
- 66. R Since dynamic routing resilience is normally required, along with policy application, eBGP routing shall be adopted as the target routing design, even where resilience is *not* deployed, such that resilience can be easily added to a non-resilient **b**-interface.
- 67. M Since eBGP peering is indirect, eBGP multi-hop support must be available and enabled.
- 68. R Given the indirect nature of the eBGP peering, loopback addresses are required for the eBGP peering establishment (and a coherent loopback address range will simplify policy).
- 69. M Given the multi-hop nature, application of the Generalised TTL Security Mechanism (RFC5082) must be modified in consideration of the hop count (see security section later within this annex).
- 70. M Traffic Selectors at the Tier-1**b** and Tier-2 Security Gateways must allow eBGP traffic between the eBGP peering loopbacks.
- 71. M Routing must be in place at both the *Neutral Host* router and *Operator* router to direct eBGP traffic (i.e. destined for the far-end peering loopbacks) towards the near-end SecGW, such that it can be transported within the **b**-interface IPsec tunnel.

72. I A minimal number¹⁹ of prefixes will be shared via eBGP in either direction (it can be considered that aggregate address ranges of S1/NG endpoints are shared) in order to keep routing scale and policy application manageable.

4.3 b-interface Routing and Resilience

73. I To provide *b*-interface resilience (independently of Tier-1*f* SecGW or Aggregation Function resilience), it is necessary to add an independent ‘B-pair’ of routers (with the hosting of these requiring geo-resilient sites in both the **Neutral Host Domain** and **Operator Domain** at a point of scalability decided by each *Operator*), with separately established IPsec tunnels and eBGP peering sessions, such that routing can be shared between the **Neutral Host Domain** and the **Operator Domain** via *both* eBGP peering sessions.

74. I An overview of resilient *b*-interface eBGP peering is illustrated in *Figure 4-5*.

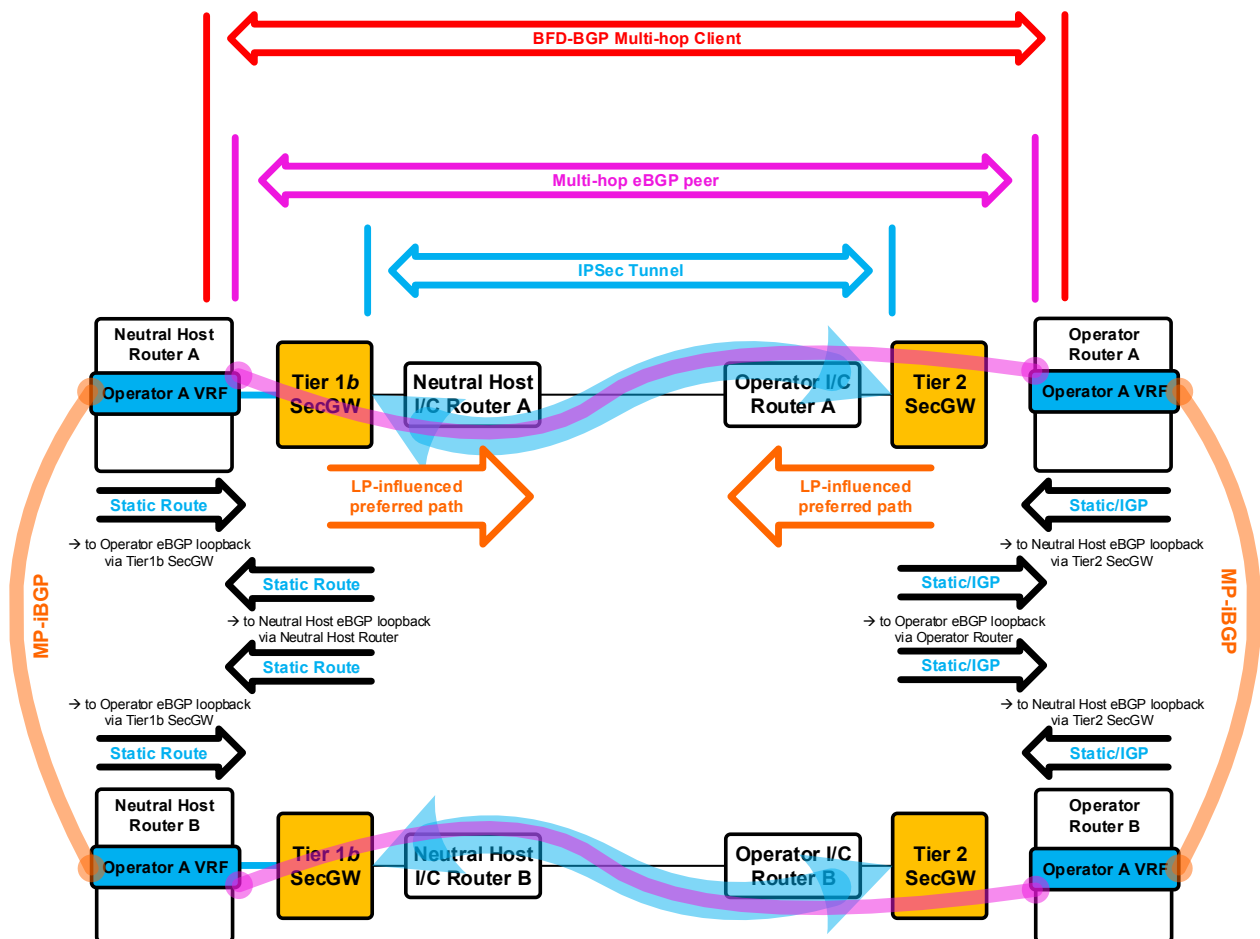


Figure 4-5 – Resilient *b*-interface eBGP Peering Approach.

¹⁹ The number of prefixes to be shared will form part of the dimensioning criteria of the overall solution.

75. I The *Neutral Host* Interconnect (I/C) routers and *Operator* Interconnect (I/C) routers are shown as dedicated devices in both the **Neutral Host Domain** and **Operator Domain**.
76. M If the SecGWs are served by separate *logical* Interconnect (I/C) routers implemented on a shared device, then the only valid routable path must be *via* the Tier1**b** to Tier2 SecGW path and must not circumvent it.
77. M It must be ensured that the eBGP peering session for each ‘side’²⁰ cannot establish, via an alternative path, routes to the loopback addresses of the other ‘side’ (since this will compromise the effectiveness of the intended path failover mechanism).
78. R It is recommended that loopback addresses are *not* advertised into the IGP at either end, but if they are required to be advertised into the IGP, correct application of IPSec Traffic Selectors (such that only the peering valid for that ‘side’ is admitted to the relevant tunnel) will ensure proper setup of the peering sessions.
79. M BFD-BGP Multi-hop Client support must be provided as a capability to support the *optional* use of BFD for the eBGP peering session.
80. I BFD for the eBGP peering session is not mandated, however it must be noted that in all cases the BGP timers may need optimising, by taking into account *Operator* and *Neutral Host* IGP settings, in order to achieve desirable failover times²¹.
81. M In order to provide deterministic routing, BGP policy will be applied to promote (from default 100) LP (Local Preference) in *both* directions for prefixes received over the A-side peering and to demote (from default 100) LP in *both* directions for prefixes received over the B-side peering.
82. I MP-iBGP (for VPN-IPv4 address family) peering between the A-side and B-side PE VRFs will clearly share the prefixes received over both peerings. Should the primary eBGP peering session drop, the routes will be withdrawn leaving the best paths those received with worse LP via the B-side peering.
83. M Resilience of functions within the **Neutral Host Domain** is required to avoid Single-Point-of-Failure risk. Resilience of the Tier-1**f** and Tier-1**b** functions are already mandated, but it must be ensured that the supporting routing and switching infrastructure is also resilient and tolerant of failures, both within the **Neutral Host Domain** and the **Operator Domain**. Resilience of other functions, such as the Aggregation Function will typically be achieved by at least local (intra-site) clustering. Where necessary, to avoid service impact on failure, geo-resilience for these functions will be deployed with

²⁰ Here referring to the routers/switches/SecGW on each side (A-side and B-side) of the resilient path.

²¹ Given that resilient paths can be expected to exist within the core networks (themselves re-converging in failure conditions), care must be taken to ensure that aggressive timers against the eBGP session (between **Neutral Host Domain** and **Operator Domain**) do not result in repeated re-convergence in failure conditions.

similar routing mechanisms as described for the *b*-interface employed within the **Neutral Host Domain** to provide resilient reachability for these functions.

84.1 The overall routing path/mechanism is illustrated in *Figure 4-6*.

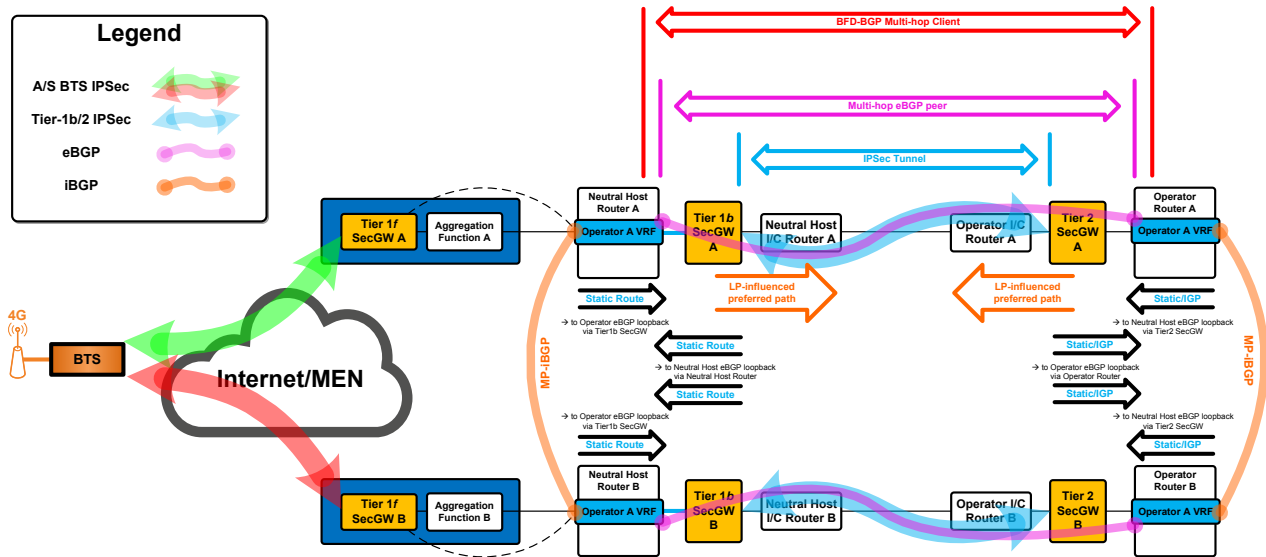


Figure 4-6 – End-to-End Overview.

85.1 The normal running traffic path is illustrated in *Figure 4-7*.

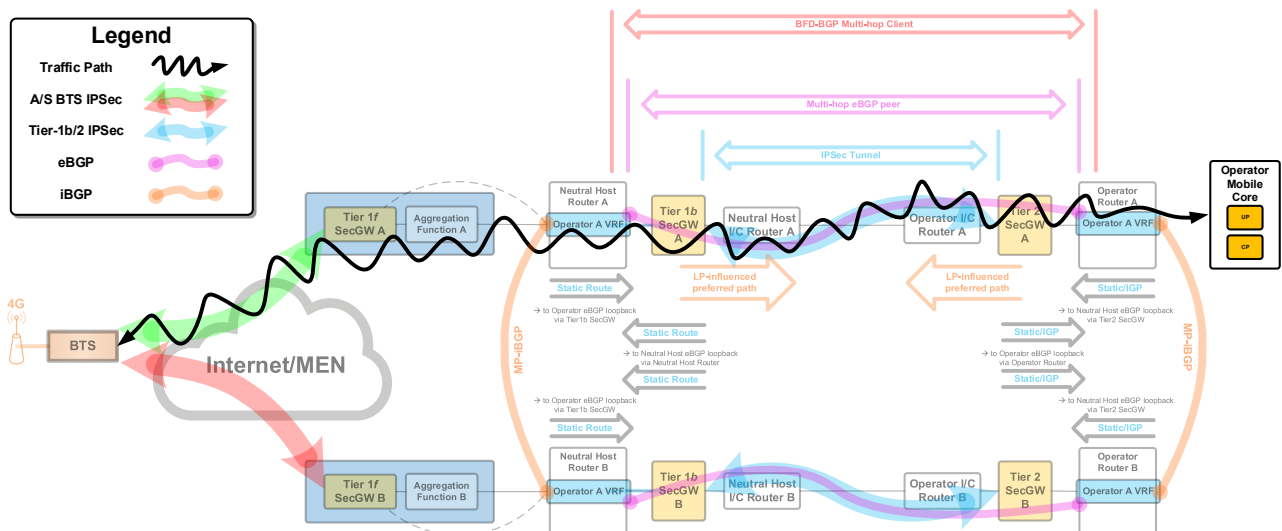


Figure 4-7 – End-to-End Traffic Path – Normal Running.

86. I Failure of any of the components shown in *Figure 4-8* will result in eBGP session failure on the A-side, resulting in the B-side learned paths becoming preferred and installed in the RIB. It should be noted that the traffic path for the *f*-interface will be unchanged in this condition – only the *b*-interface routing path has changed.

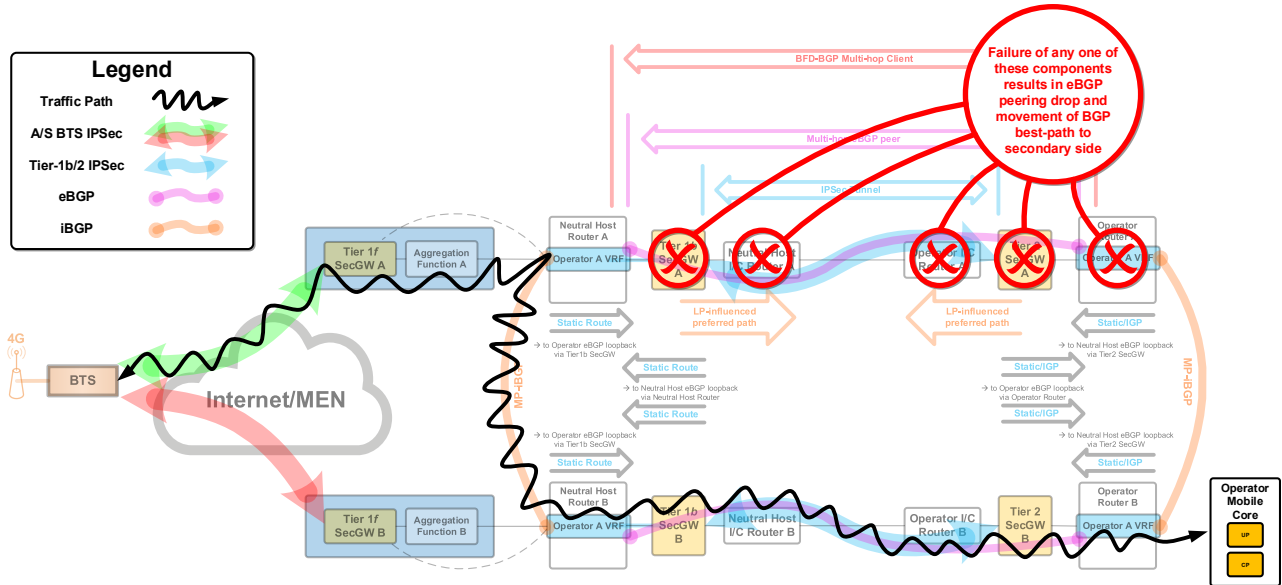


Figure 4-8 – End-to-End Traffic Path – A-side eBGP Session Down.

87. I Failure of any of the components shown in *Figure 4-9* will result in loss of connectivity for the *f*-interface. Where supported by the BTS, failover to secondary site Tier-1*f* SecGW and Aggregation Function will occur (noting that in this model they do not share state with the A-side).

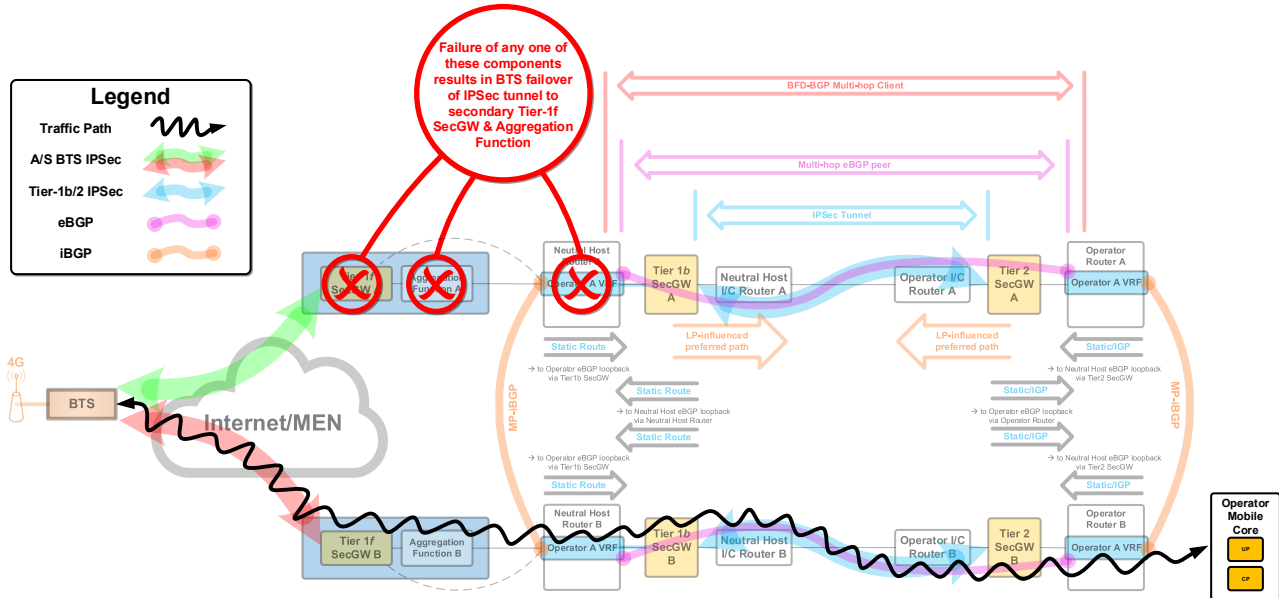


Figure 4-9 – End-to-End Traffic Path – A-side Tier-1*f*/HeNB-GW/NH-Router Failure.

88. I *Figure 4-9* shows the *b*-interface routing path has also moved to the B-side – this is as a result of failure of the **Neutral Host Domain** A-side router. It should be noted that *f*-interface and *b*-interface resilience are independent. If solely either the A-side Tier-1*f* SecGW or Aggregation Function had failed, the *f*-interface would move to the B-side Tier-1*f* SecGW and Aggregation Function, but onward *b*-interface transport would follow the (still) preferred A-side path learned from the A-side via MP-iBGP.
89. I Advertisement of routing information via the eBGP sessions between **Neutral Host Domain** and **Operator Domains** is independent from the steering of traffic into specific IPsec tunnels, which is achieved by the Traffic Selector definition. The options for separation of traffic types into tunnels are described in *Section 7.2*.

4.4 IP Addressing

90. I Tunnel-outer addressing for establishment of *f*-interface IPsec tunnels are outside of the *Operator* routing domain, require no reachability from the **Operator Domain** and are therefore entirely the responsibility of the **Neutral Host Domain**.
91. R Tunnel-outer addressing for establishment of *b*-interface IPsec tunnels requires an agreement of appropriate addressing between the **Neutral Host Domain** and each **Operator Domain**. Given that

this addressing is solely to establish the **b**-interface tunnel, it is not necessary for it to be routed from within other *Operator* VRFs and therefore is not expected to cause an issue in allocation.

92. M Tunnel-inner addressing relating to BTS endpoints must be **unique** per-*Operator*.
93. M Per-*Operator* unique²² IPv4 addressing will be required as a minimum for tunnel-inner address ranges relating to BTS endpoints.
94. R Where an Aggregation Function is deployed, providing both S1-AP and S1-U aggregation or NG aggregation, it is possible and recommended to utilize a private network managed by the **Neutral Host Domain** from that Aggregation Function towards the access (i.e. BTS nodes).
95. I When an Aggregation Function is deployed, it is expected that only a small number of aggregated S1/NG endpoints need to be uniquely²² addressed from the **Operator Domain**, which should significantly ease achievement of the IP-uniqueness requirement.
96. M Tunnel-inner addressing relating to mobile core endpoints must be **unique** per-*Operator*.
97. M Per-*Operator* unique²² IPv4 addressing will be required for mobile core addressing, such that traffic from the **Neutral Host Domain** can reach the appropriate *Operator*-specific VRF and such that the correct **b**-interface IPsec tunnel can be selected based on unique Traffic Selector content.
98. R In order to accommodate evolution of *Operator* RAN environments towards IPv6 address space, along with the potential of IPv6 to significantly simplify IP uniqueness requirements across multi-*Operator* solutions, it is strongly advised that IPv6 capability of selected equipment is properly considered (along with any relevant limitations associated with IPv6 implementations), even where not required in initial deployments.
99. R It is recommended that the evolution towards IPv6 addressing follows a 3-stage process:

Stage 1 - IPv6 User Plane addressing, followed by,
Stage 2 - IPv6 Inner addressing, followed by,
Stage 3 - IPv6 Outer addressing,

where each stage builds on the developments introduced in the previous stages. It should be assumed that each *Operator* will evolve through the three stages at different rates, meaning the *Neutral Host* will be required to support IPv4 addressing and all three stages of IPv6 evolution concurrently.

²² It should be noted that future evolution towards a vRAN model, or deployment of per-*Operator* logical instances for the end-to-end solution, will present the opportunity for IP addressing overlap to be achieved.

Stage 1 – IPv6 User Plane addressing

For UE addressing towards an IPv6 enabled mobile core, IPv6 from the UE will continue to be encapsulated within GTP at the BTS. Initially the BTS 'inner' addressing associated with the GTP Tunnel Endpoint Identities for S1-U/N3 and SCTP paths for S1-C/N1/N2 can remain IPv4 (noting that, at this point, uniqueness of IP Inner addressing across *Operators* is not solved).

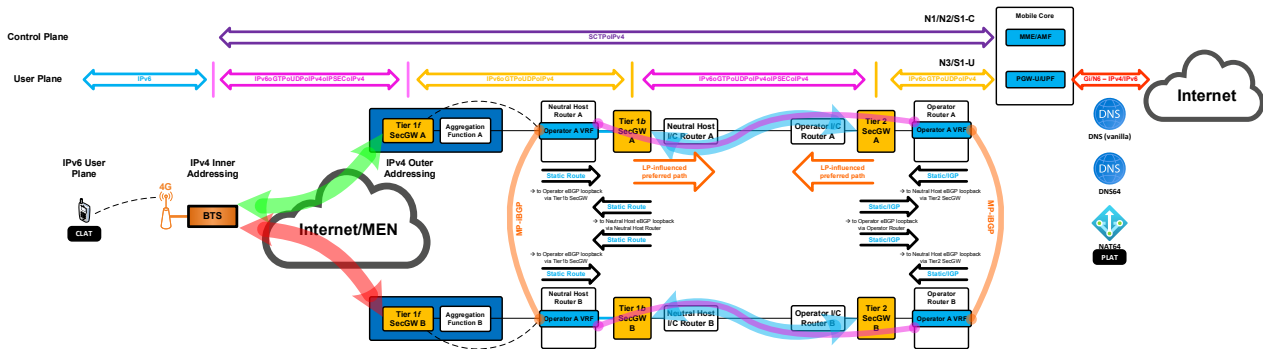


Figure 4-10 – IPv6 evolution: Stage 1 – IPv6 User Plane addressing

Stage 2 – IPv6 Inner Addressing

Evolution of the BTS 'inner' addressing to IPv6 will allow uniqueness across *Operators*. It will simplify both the addressing considerations within the **Neutral Host Domain** and onward scalability. It is assumed that IPv6 tunneling in IPsec/IPv4²³ 'outers' is supported by the SecGW functions and, as such, it should not be necessary to alter the tunnel 'outer' addressing (for any IPsec span) when using IPv6 Tunnel End Points. It must be noted however, that at this point both the **Neutral Host Domain** and **Operator Domain** must route IPv6 traffic in their cores (since it will no longer be encapsulated within IPv4).

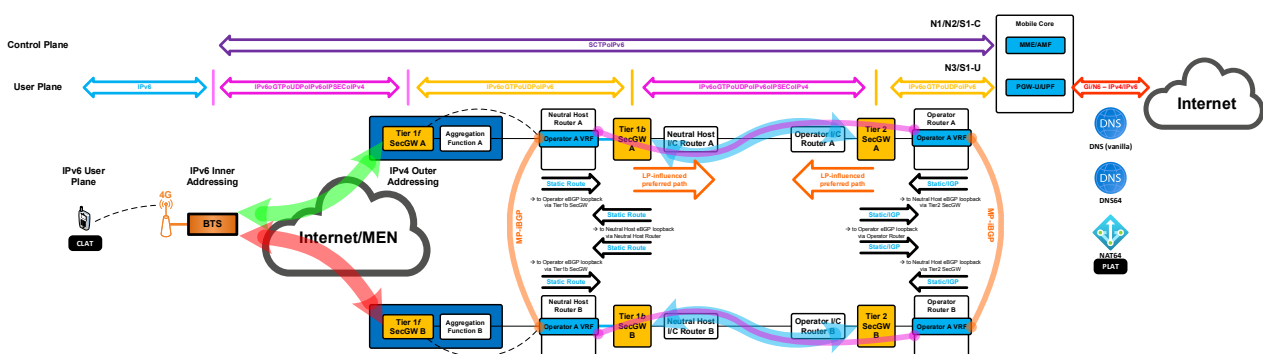


Figure 4-11 – IPv6 evolution: Stage 2 – IPv6 Inner Addressing

²³ RFC: 4891, Using IPsec to Secure IPv6-in-IPv4 Tunnels

Stage 3 – IPv6 Outer Addressing

Evolution of the BTS 'outer' addressing to IPv6 removes the requirement for IPv4 support in the transport path altogether.

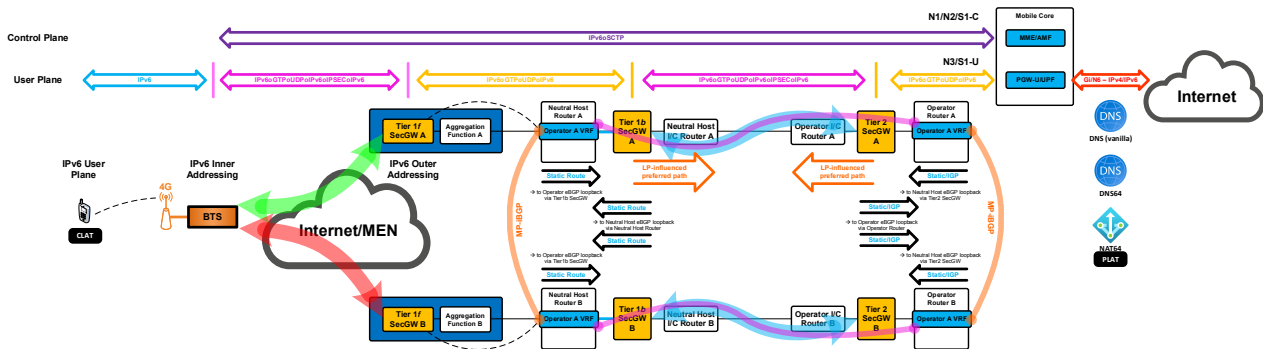


Figure 4-12 – IPv6 evolution: Stage 3 – IPv6 Outer Addressing

- 100. I Any functions deployed within the **Neutral Host Domain** (e.g. aggregation/steering functions) will need to support IPv6 traffic from (and including) Stage 2 – IPv6 Inner Addressing.
- 101. I Any decision around UE single-stack (IPv4 only) or dual-stack (IPv4/IPv6) working resides separately with each *Operator*.
- 102. R The *Operator* is responsible for the provision of appropriate techniques, such as DNS64, NAT64 and 464XLAT PLAT, to support connectivity of IPv6 addressed UEs to IPv4 addressed services within either the **Operator Domain** or within external networks (e.g. Internet).
- 103. R The small cell BTS components are to be presented IP addressing for tunnel-outers by means of DHCP.
- 104. R It is required that DHCP server provision (with appropriate per-venue pools) be provided either at the venue (either as a function of the CPE or separately) for Internet connected venues, whilst DHCP server provision (and appropriate pool configuration) could be made either at the venue or **Neutral Host Domain** for privately connected venues.
- 105. I DHCP VSOs or other custom DHCP options may be required dependant on the deployed BTS solution. Determination of these requirements (and provision of appropriate configuration) is the responsibility of the *Neutral Host*.
- 106. M Addressing schemes must be structured, such that BTS nodes can be constrained to specific per-venue addressing pools. This is expected to be an essential component to the location lock capability in order to help determine whether a known device (e.g. by MAC address) has been allocated an IP address from a different address pool (and has therefore been moved outside of the scope of its original venue). Capability within the **Neutral Host Domain** to identify such changes will be required.

4.5 DNS/FQDN

107. I An overview of DNS resolution requirements (and responsibility) is illustrated in *Figure 4-13*:

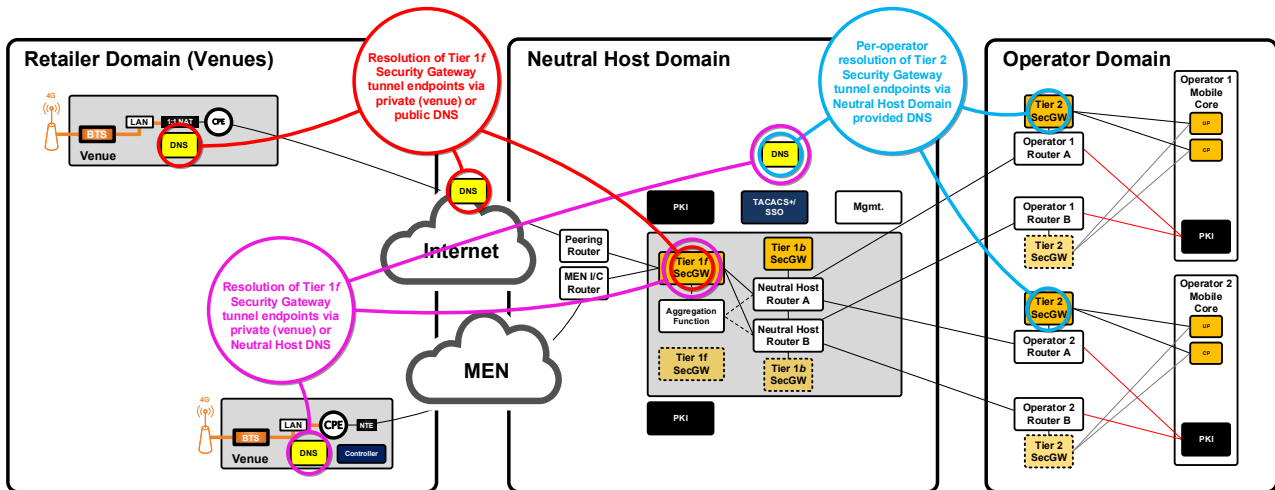


Figure 4-13 – DNS Resolution Overview.

108. M In order to simplify BTS configuration, along with (potentially) certificate generation, DNS resolution of FQDNs is mandated in order to resolve IP addresses for tunnel endpoint targets in the *f*-interface.
109. I DNS resolution of FQDNs not only simplifies initial provisioning, but also simplifies future platform scale and evolution, where migration to different tunnel endpoints may be desirable (and is otherwise cumbersome to achieve).
110. M In order to improve the DoS/DDoS ‘attack surface’ characteristics, multiple public addresses must be provided on the Tier-1*f* SecGW *f*-interface Internet-facing context as tunnel endpoints.
111. I A single public IP address for tunnel establishment into this untrusted Internet context could result in total service loss for Internet connected venues for all *Operators* with no blocking-mitigation option, should that single tunnel endpoint address come under attack.
112. R FQDNs for publicly addressed tunnel endpoints should be populated in public DNS and propagated to the wider DNS hierarchy in the Internet, such that tunnel endpoint addresses can be resolved by use of FQDN, rather than manually configured IP addressing (and likewise, certificates can be FQDN based).

-
113. R Distinct A-records for the set of tunnel endpoints should be created, with a further round-robin FQDN C-NAME (i.e. a C-NAME with round-robin enabled) resolving to those A-records.
114. I Using a C-NAME, with round-robin enabled, partially mitigates an attack against a single tunnel endpoint address. An attack may result in an established IPsec tunnel to the affected public tunnel endpoint address dropping. But as FQDN resolution will again be used in order to attempt re-establishment, retry will occur until such point as a working (and not blocked) tunnel endpoint address is provided by the C-NAME query.
115. R Since the FQDN records will be present in a public DNS, it is recommended that they are both meaningful *and* to some extent obfuscated, such that they do not become an obvious target to an attacker. FQDN allocation will require review and approval by *Operator* security policy functions.
116. I It should be noted that DNS resolution against these A-records and C-NAMEs could alternatively be provided within the venue, or from the **Neutral Host Domain**, but it is assumed as a base case that they are resolvable via public DNS lookup (and as such must in any case be propagated through the public DNS infrastructure).
117. I Where *f*-interface connectivity is provided via a MEN (or other ‘private’ connectivity), FQDN resolution must either be provided by local DNS resolution at the venue, or from the **Neutral Host Domain**, noting that resolution in the **Neutral Host Domain** will need to take place prior to IPsec tunnel establishment and therefore will require a VLAN extension specifically for the purpose of this initial DNS resolution.
118. R Given the (typically) private nature of addressing for the Tier-1**b** SecGW to Tier-2 SecGW interface, it will be necessary for DNS resolution to be provided within the **Neutral Host Domain** for this component.
119. R Connectivity between the **Neutral Host Domain** Tier-1**b** SecGW and the **Operator Domain** Tier-2 SecGW is anticipated to typically be achieved via private connectivity models, not exposed to DoS/DDoS threats. However, given that this interface *could* be achieved for small *Neutral Hosts* by means of Internet connectivity and given that FQDN use will in any case simplify any future addressing changes for this interface, it is required that FQDN resolution is used for this interface.
120. R In all cases a robust and heavily resilient DNS infrastructure is required in order to ensure resolution availability.
121. R Where UEs are running a single IPv6 stack, the *Operator* is required to support the necessary DNS techniques (notably DNS64 in conjunction with NAT64) to allow those UEs to reach IPv4 addressed internal or external services.

5 TRAFFIC AGGREGATION

122. I It is highly desirable, from a mobile core perspective, notably the Mobility Management Entity (MME) or Access and Mobility Management Function (AMF), to provide a set of aggregated S1/NG interfaces for the deployment in the **Retailer Domain** and **Neutral Host Domain**, such that scalability of the mobile core elements does not present a key limitation in the overall solution and to simplify the IP addressing uniqueness requirements across *Operators*. As such, it is required that the presentation from the **Neutral Host Domain** towards the **Operator Domain** is aggregated.
123. M Where an *Operator* requires multi-PLMN support or two or more *Operators* require support for MOSS, S1/NG aggregation becomes a mandatory requirement. The *Neutral Host* must support not only aggregation functions, but traffic steering of multiple PLMNs to distinct EPC/5G core elements/pools.²⁴
124. R Where an *Operator* requires multi-PLMN support, it is required that the *Neutral Host* solution can apply end-to-end QoS and Priority (including RRC, QCI-specific level differentiation), Pre-emption, Access Class Barring and MTPAS procedures separately per PLMN, in order to adequately (and independently) support the requirements each PLMN and/or perform locking of cells for one PLMN whilst leaving the other PLMN(s) operational.
125. R Where an *Operator* requires multi-PLMN support, it is required that the solution supports sufficient overall MME/AMF pools to support the selective steering of PLMNs to distinct mobile cores (or distinct endpoint pools), either at a single-*Operator* or multi-*Operator* level.
126. M Where an *Operator* requires multi-PLMN support, it is required that metrics and operational reporting data be separable between deployed PLMNs, as these will need to be available to typically separate operational teams within the **Operator Domain**, or in the case of MOSS deployments, must be available in distinct form to multiple *Operators*.
127. R Where an *Operator* requires multi-PLMN support, it can be assumed that support for a minimum of five PLMNs is required for the UK, but selection of products with greater PLMN scalability (or roadmapped scalability) is recommended in order to adequately accommodate potential future need.
128. M The total number of S1/NG interfaces presented towards each **Operator Domain** shall be agreed between the *Neutral Host* and the *Operator* and will be reviewed at least on an annual basis.

²⁴ The aggregation and steering functions can be combined or separated. They may be achieved at the radio solution, where appropriate capabilities exist in the *Neutral Host* selected solution and/or by means of a separate distinct (centralised) Aggregation Function. It should be noted that in the **Operator Domain**, MME/AMF may provide distinct instances/endpoints per PLMN, but equally may serve multiple PLMNs from shared instances/endpoints and therefore it is recommended that *Neutral Host* selected solutions support both models.

129. I Aggregation (especially where achieved at a high ratio) can be expected to simplify the IP uniqueness requirements (i.e. to achieve non-ambiguous routing across the **Neutral Host Domain** for each *Operator*), simply due to this Aggregation Function minimising the number of uniquely addressable endpoints required.
130. I In some deployment models, a Controller function will exist, typically at the venue, providing the first layer of BTS aggregation and simplifying a number of deployment topics such as RF planning/optimisation, BTS to BTS handover and management (at the venue level in the **Retailer Domain**). Due to the aggregation of hosted BTS nodes, the Controller node is expected to appear as a single BTS from an MME perspective (i.e. it presents an aggregated representation of the S1/NG interfaces).
131. I The Controller functions can be deployed in two models:
- **Centralised** (i.e. deployed within the **Neutral Host Domain**);
 - **Localised** (i.e. deployed within the venue component of the **Retailer Domain** – the expected typical deployment).
- Where **centralised**, the Controller node will typically host a number of distinct **Retailer Domain** venues (with this model typically being appropriate to aggregate traffic from multiple small venues where it is not commercially viable to deploy a **localised** Controller node).
132. M Where the Controller function is not deployed as an initial tier of S1/NG aggregation, the requirement to present aggregated interfaces to the *Operators* mandates that a dedicated Aggregation Function is provided within the **Neutral Host Domain**.
133. M It is required that S1-AP/S1-U/NG interfaces are aggregated by the Aggregation Function, such that other treatment of traffic flows (notably routing and security policy) does not have to be applied differently for the S1-AP/S1-U/NG interfaces (and such that all are presented in an equally aggregated manner towards the **Operator Domain**).
134. I Where IPv6 capability is deployed, it will be necessary for both the routing domain and any aggregation/steering functions deployed in the **Neutral Host Domain** to include IPv6 support.

6 DISAGGREGATED RAN

135. I The architecting principles embedded within the Neutral Host In-Building (NHIB) solution can be extended to incorporate disaggregated RAN concepts for the purposes of supporting 5G and evolving 4G architectures (notably ORAN).

By design NHIB is separated into operational domains of responsibility and into each of these disaggregated RAN components can be deployed. However, the mapping of disaggregated RAN components into domains must take into account practical operational and technical capabilities and therefore the mapping is not arbitrary.

In this version of the specification we consider the first step towards RAN disaggregation applied to NHIB. Later versions of this specification may extend the disaggregation options further as radio vendor, *Retailer*, *Neutral Host*, *Operator*, front-haul and mid-haul capabilities grow within an established NHIB ecosystem.

136. I In this specification RAN disaggregation is limited solely to the **Neutral Host Domain** and the **Retailer Domain** (i.e. the **Operator Domain** is not in scope). The architecture of the disaggregated RAN option is shown in

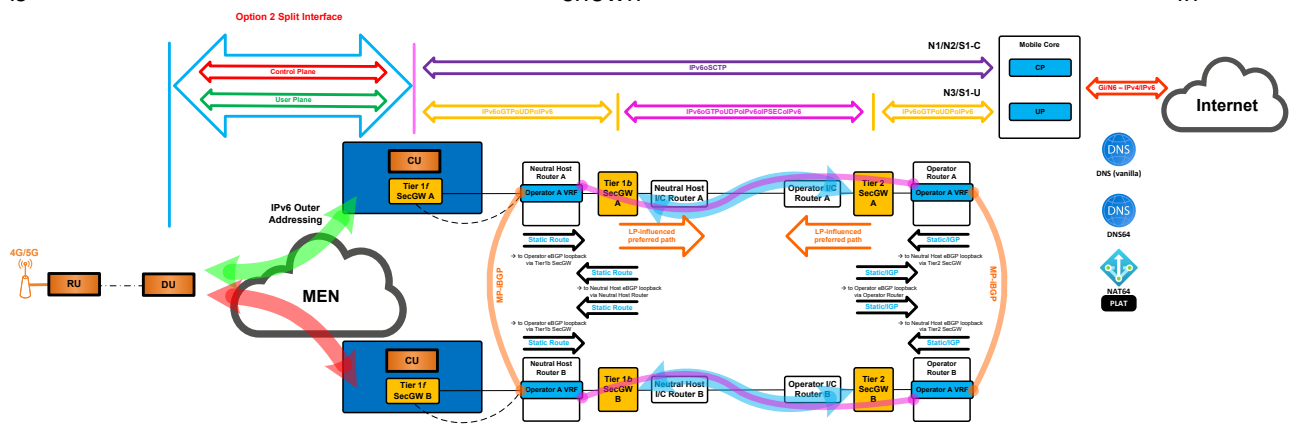


Figure 6-1.

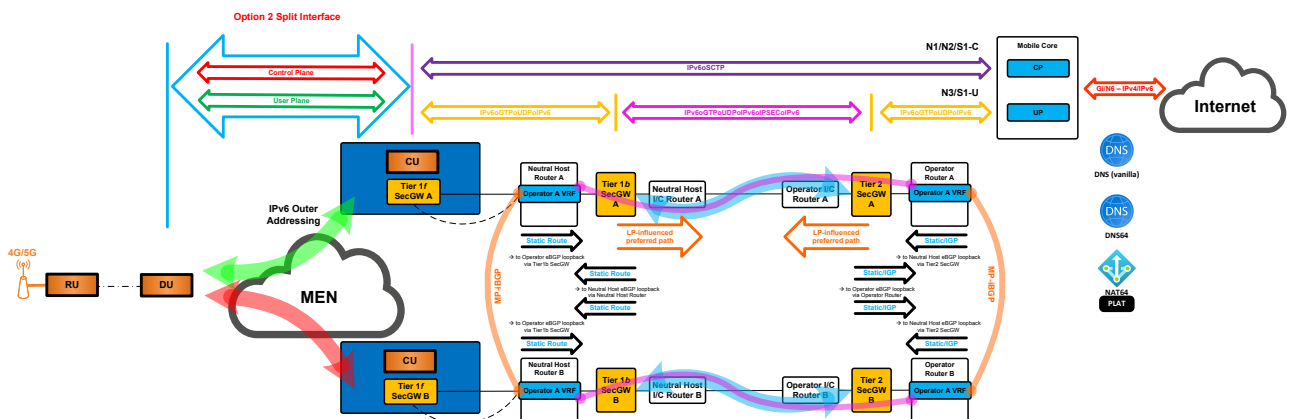


Figure 6-1 – Disaggregated RAN.

The proposed first step towards disaggregated RAN shows the CU residing in the **Neutral Host Domain** (not the **Operator Domain**). Thus, the *Neutral Host* maintains full control over the operation of the radio solution which is comprised of CU/DU/RU components.

The RU component resides in the **Retailer Domain** (i.e. at the venue).

The DU component would most likely reside in the **Retailer Domain**, but it could alternatively be deployed in the **Neutral Host Domain**. Its location would depend to a greater extent on the availability and cost of suitable front-haul (RU to DU) and mid-haul (DU to CU) connectivity.

137. I Locating the CU in the **Neutral Host Domain** delivers a number of advantages:

- Back-haul connectivity from the CU towards the *Operator's* core (via the *b*-interface) remains unaltered in terms of its routing, resilience, security, QoS and performance requirements. This presents a simpler first RAN disaggregation step;
- Similarly, aggregation and steering functions within the **Neutral Host Domain** can (and should) equally apply to disaggregated RAN traffic;
- It enables the *Neutral Host* and *Operator* to deploy simpler inter-CU routing rules and policies when implementing X2/Xn interfaces between potentially multiple *Neutral Hosts* and multiple *Operators*. The *Neutral Host* simply routes traffic from their CU towards the required IP destination. The *Operator* is not required to trombone X2/Xn traffic back towards a *Neutral Host* or other *Operator* and is therefore only required to terminate X2/Xn traffic within its own domain;
- The *Neutral Host* can maintain consistency of vendor during the first step of RAN disaggregation (i.e. DU/CU, potentially RU, from the same vendor). Later the *Neutral Host* can look to inter-operate with other vendor kit as the technology matures and becomes vendor agnostic (the ORAN vision);
- The *Neutral Host* avoids a multiplicity of potentially complex inter-operability and security tests which would be required if separate and different (vendor) CU were hosted across multiple **Operator Domains**;
- The *Neutral Host* maintains the responsibility to operate and manage the radio solution on behalf of the *Operator* (a key feature of NHIB concept);
- The *Neutral Host* maintains control over the synchronization of the overall radio solution.

138. R It is required that a disaggregated solution aligns to 3GPP²⁵, in that:

- F1-C interface shall support confidentiality, integrity and replay protection;
- All management traffic carried over the CU-DU link shall be integrity, confidentiality and replay protected;
- The gNB shall support confidentiality, integrity and replay protection on the gNB DU-CU F1-U interface for user plane;
- F1-C and management traffic carried over the CU-DU link shall be protected independently from F1-U traffic.

Noting that the above requirements allow to have F1-U protected differently (including turning integrity and/or encryption off or on for F1-U) from all other traffic on the CU-DU (e.g. the traffic over F1-C).

²⁵ 3GPP TS33.501 *Security architecture and procedures for 5G System*, V16.9.1, section 5.3.9.

139. R It is recommended that the *Neutral Host* maintains the capability to further encrypt the F1-U interface using IPsec, however since this may already be an encrypted user-plane traffic flow, the hosted *Operator* may declare and agree that additional IPsec encryption is not required on the F1-U interface.
140. R Where IPsec is (optionally) used for the purposes of authorisation and demarcation on the F1-U interface, the hosted *Operator* may declare and agree that ESP NULL encryption is applied to the Child-SA associated with the F1-U traffic in order to minimise compute resources (cost and power).
141. I The omission of IPsec from the F1-U interface altogether, and therefore the avoidance of SecGW in the F1-U path, can be expected to allow further optimization of traffic paths for F1-U, potentially leading to greater coverage reach of a CU (as *Neutral Host* datacentre routing diversity and routing options grows).

7 SECURITY

7.1 Security Domains

142. I An overview of the RED, AMBER and GREEN security domains and security gateway components is illustrated in *Figure 7-1*, with respect to the trustiness/vulnerability of the NHIB system in terms of mobile network security as far as the Operator is concerned.

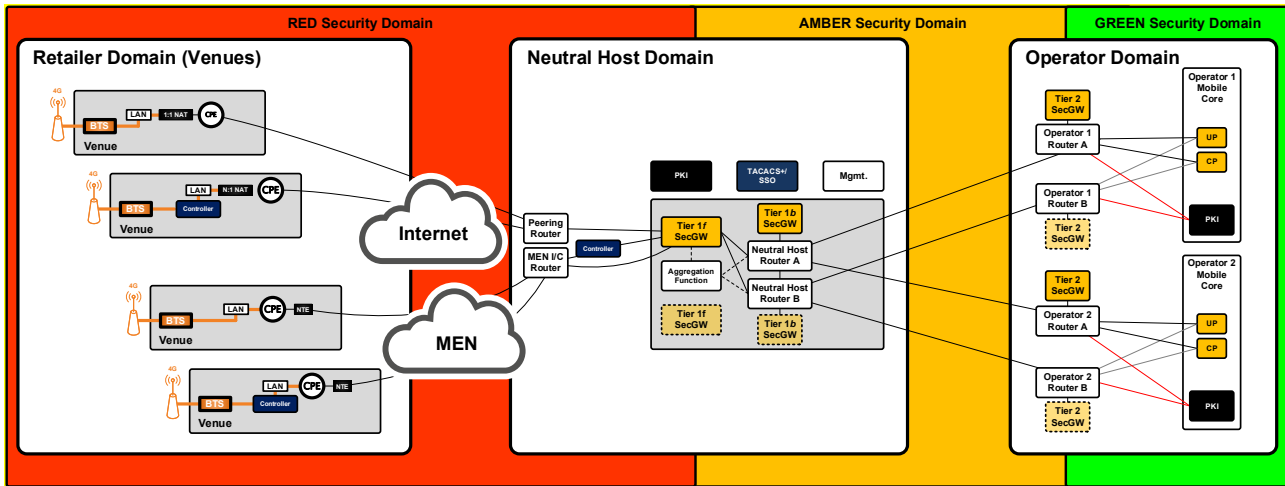


Figure 7-1 – Security Domain Overview.

143. M The security requirements within each security domain is set out in *Table 7-1*:

Table 7-1 – Security Domain Requirements

Security Domain Requirements		
RED	AMBER ²⁶	GREEN
Untrusted	Formal physical controls (section 7.3, 7.6)	As per Operator GREEN definitions.
	Role based formal logical access controls (RBAC/SSO) (section 7.6, 9)	
	Access logging (section 7.6)	
	Configuration logging (section 7.3)	
	Auditing including interval-based penetration testing	
	Multi-tenant management separation (section 9)	
	Multi-tenant reporting separation (e.g. SNMP and syslog) (section 9)	
	Minimisation of clear traffic domain in line with NHIB specification. (section 7.3, 7.6)	
	Protection of cleartext traffic scope (e.g. logical or physical controls e.g. armoured conduit) (section 7.3, 7.6)	
	Logical separation of per-Operator traffic wherever possible in line with NHIB specification (section 4, 5, 7)	
Approval of security controls by Operator Domains (section 7)		

144. R A first tier of SecGW connectivity facing towards the **Retailer Domain**, referred to as a Tier-1f SecGW, is required within the **Neutral Host Domain**. The Tier-1f SecGW terminates per BTS or per Controller

²⁶ Further details relating to each of the security domain requirements can be found by reference to ISO 27001 and TSCoP (see footnote 7).

IPsec tunnels within the RED security domain. While terminating in an untrusted (RED) domain, which is exposed to the general public, the IPsec tunnel endpoints are themselves trusted.

- 145. R A second tier of SecGW connectivity facing towards the **Operator Domain**, referred to as a Tier-1**b** SecGW (governed by **Operator Domain** policy), is required within the **Neutral Host Domain**. The Tier-1**b** SecGW serving the **b**-interface connection terminates IPsec tunnels within the AMBER security domain.
- 146. R A Tier-2 SecGW terminating the IPsec tunnels is required within the **Operator Domain** to terminate the **b**-interface connection from the **Neutral Host Domain**. The Tier-2 SecGW terminates IPsec tunnels at the edge of the AMBER security domain.
- 147. I The *Operator* core facing side of the Tier-2 SecGW (which resides within the **Operator Domain**) terminates interfaces towards the GREEN security domain.
- 148. M The Tier-1**b** SecGW function in the **Neutral Host Domain**, presenting **b**-interface connectivity towards the multi-*Operator* cores, must be at least logically separate from the Tier-1**f** SecGW used for the **f**-interface terminations.
- 149. R BTS within the **Retailer Domain** obtain their trusted status through compliance with the principles set out in TS 33.320²⁷.

7.2 IP Security (IPSec)

- 150. R The **f**-interface component of IPSec connectivity will be the responsibility of the *Neutral Host*, in compliance with the minimum requirements stated in the Security Parameters sub-annex²⁸, which will be agreed between *Operators* and provided as part of the *Neutral Host* engagement process.
- 151. R Given that the **f**-interface interface will be presented for some venues via public Internet, with others being presented via Metro Ethernet (i.e. private) services, it is required that logically separate contexts be created on the Tier-1**f** SecGW with *Operator* separation where appropriate, referred to as the **f-interface-internet-context** and **f-interface-private-context**. Appropriate policy should be in place to avoid any interconnection between these contexts.

²⁷ 3GPP TS 33.320 Security of Home Node B (HNB) / Home evolved Node B (HeNB).

²⁸ JOTS NHIB Specification – Annex 1 – Security Parameters (most up to date version applies).

152. I An overview of the **b**-interface IPsec connectivity approach is illustrated in *Figure 7-2*.

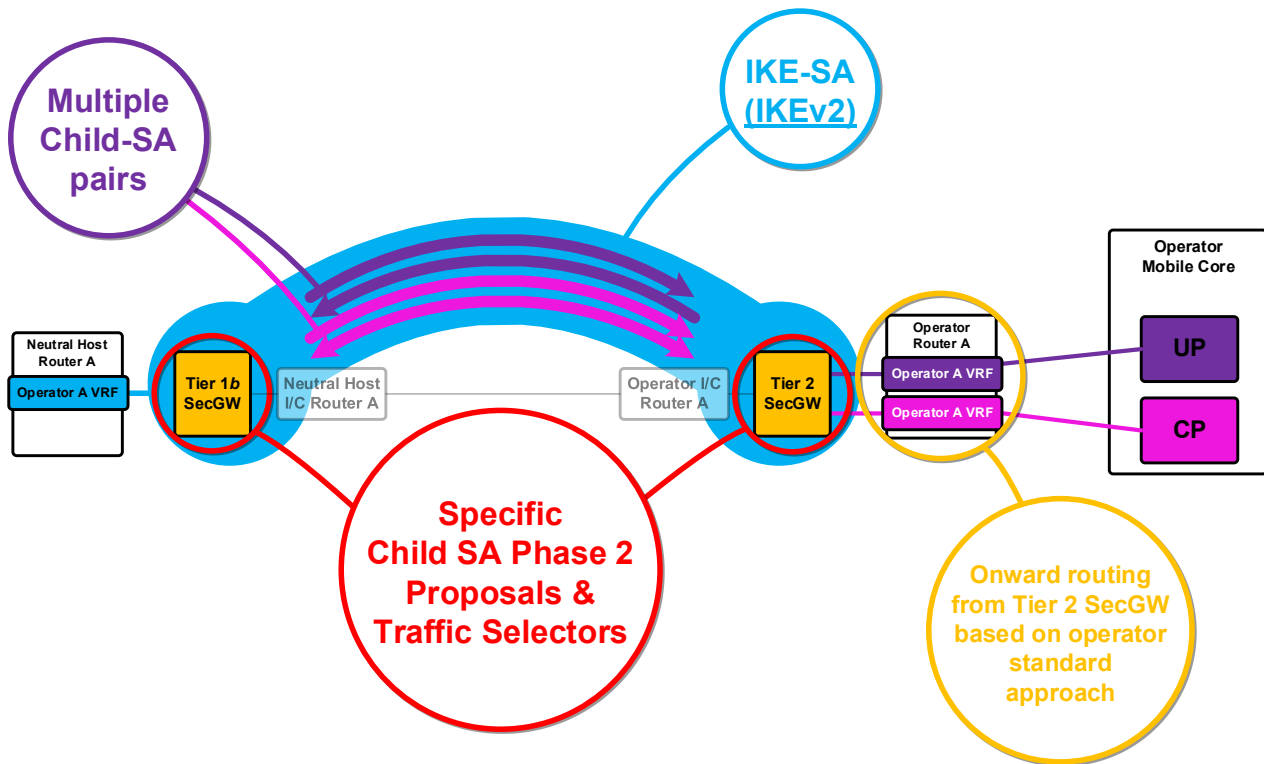


Figure 7-2 – b-interface IPsec Connectivity.

153. M **b**-interface IPsec must meet the minimum requirements stated in the Security Parameters²⁸ sub-annex, which will be agreed between *Operators* and provided as part of the *Neutral Host* engagement process.
154. I The **b**-interface component of IPsec connectivity consists of a set of IPsec tunnels established from the Tier-1**b** SecGW in the **Neutral Host Domain** to the set of Tier-2 SecGWs in the **Operator Domain**.
155. R These IPsec tunnels will be strictly based on tunnel selection and traffic forwarding by use of defined Traffic Selectors of more than one distinct IP source/destination pair with appropriate port/protocol information (i.e. from multiple BTS and/or Controller sources that may not be contiguous and towards multiple MME and SGW address ranges on the *Operator* side), and it will be necessary (initially) for these pairs to be achieved by establishment of multiple Child-SAs under a single IKE-SA (IKEv2) association²⁹.
156. R In order to achieve further scalability (either for maximum-traffic-per-tunnel constraints, maximum traffic per Child-SA constraints or due to Child-SA per IKE-SA scale limitations), there is a requirement

²⁹ Since this will require a single IKE-SA pair and only a limited number of Child-SAs per IKE-SA for each *Operator*, it is not anticipated that this approach will present scalability limitations, although max-traffic-per-tunnel limitations should be considered.

to be able to establish multiple IKE-SAs (i.e. Phase 1s), each with a subset of (Phase 2) Child-SAs (to cover the overall footprint), such that traffic can be spread across multiple tunnels and multiple Child-SAs per tunnel for the Tier-1**b** to Tier-2 SecGW path³⁰.

- 157. I Given that some IPSec implementations will not allow multiple IKE-SAs between the same tunnel endpoints, it *may* be necessary to present multiple tunnel endpoint addresses on the SecGW functions to fulfil this purpose.
- 158. I If required, it will also be possible for multiple eBGP sessions to be run (across potentially disparate SecGW devices), with policy on each carrying a subset of the overall routing information (and with the appropriate SecGW nodes having aligned Traffic Selectors), but this is not typically expected and adds complexity. This should therefore only be considered an option for scale and resilience if required in future³¹.
- 159. R The **b**-interface IPSec tunnel will be initiated only *from* the **Neutral Host Domain** *towards* the **Operator Domain** and configuration will be required to enforce this.
- 160. R IPSec tunnel configuration will require seamless (i.e. make-before-break) re-authentication, with appropriately short timers to be used, in order to force re-assessment of the certificate regularly (and therefore invoke any revocation as a result of addition to the PKI CRL).
- 161. I A number of valid approaches for separating traffic into discrete tunnels exist, with the simplest being presentation of all traffic types as different Child-SAs within the same Phase 1 – see *Figure 7-3*.

³⁰ How this is achieved technically, depends on the capabilities of the SecGW technical solution employed. For example, in some implementations the creation of multiple tunnels between identical IPSec peer addresses is possible, whilst in others different IPSec peer addresses are required for the creation of multiple tunnels.

³¹ Other approaches, such as multiple eBGP sessions between the **Neutral Host Domain** and **Operator Domain** with equal cost hashing via eBGP multipath and policy routing into tunnels at the Tier-1**f**/Tier-2 devices, may be more practical to achieve scale but not all vendor capabilities will support this approach.

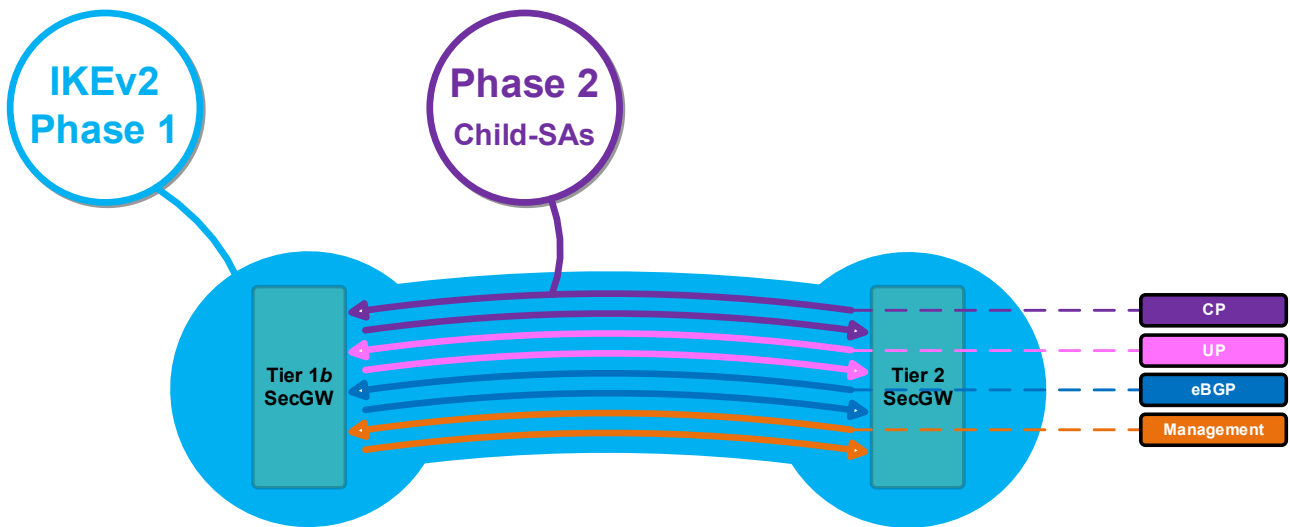


Figure 7-3 – Single IPsec Phase 1 – All traffic types in single tunnel.

162. I Where further separation is used, it must be noted that since the eBGP sessions carry routing information for *all* traffic flows (independently from tunnel selection by Traffic Selector definition), appropriate consideration should be made regarding allocation of traffic types into tunnels such that loss of forwarding plane tunnels (i.e. the S1/NG bearers) does not also result in loss of management or routing information supporting that management function.
163. I Options for traffic separation across multiple tunnels is shown in *Figure 7-4* and *Figure 7-5* (noting that the S1/NG tunnel replication model can be used across single or multiple SecGW instances at either side in order to achieve traffic scalability and that separation could be achieved by specific distinct traffic selectors or multipath hashing with policy based tunnel routing).

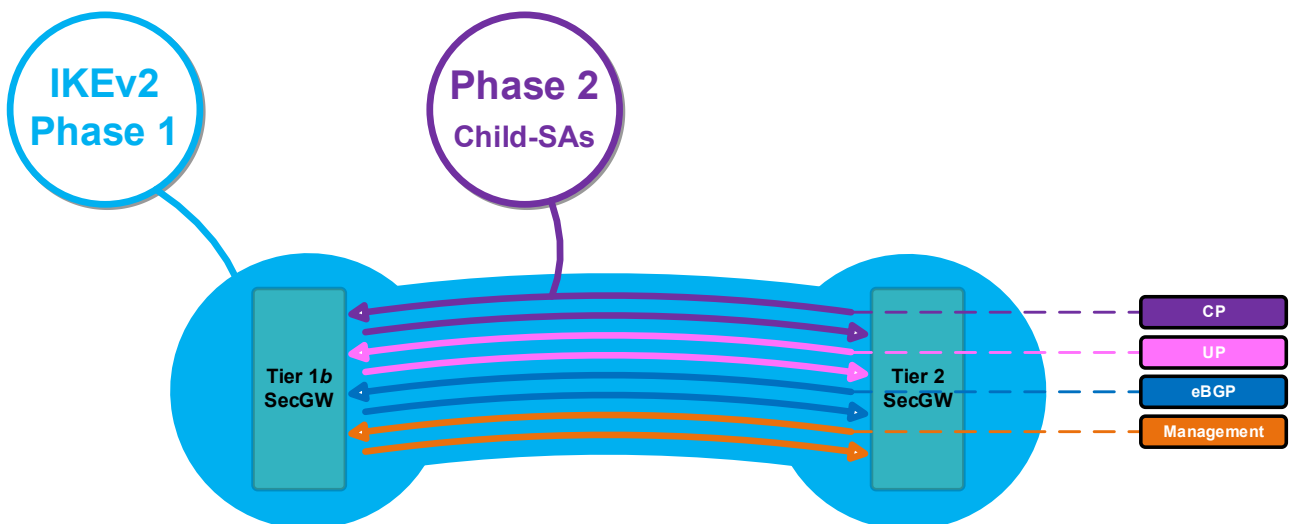


Figure 7-4 – Dual IPsec Phase 1 – Control/Management Plane and S1/NG separation across two tunnels.

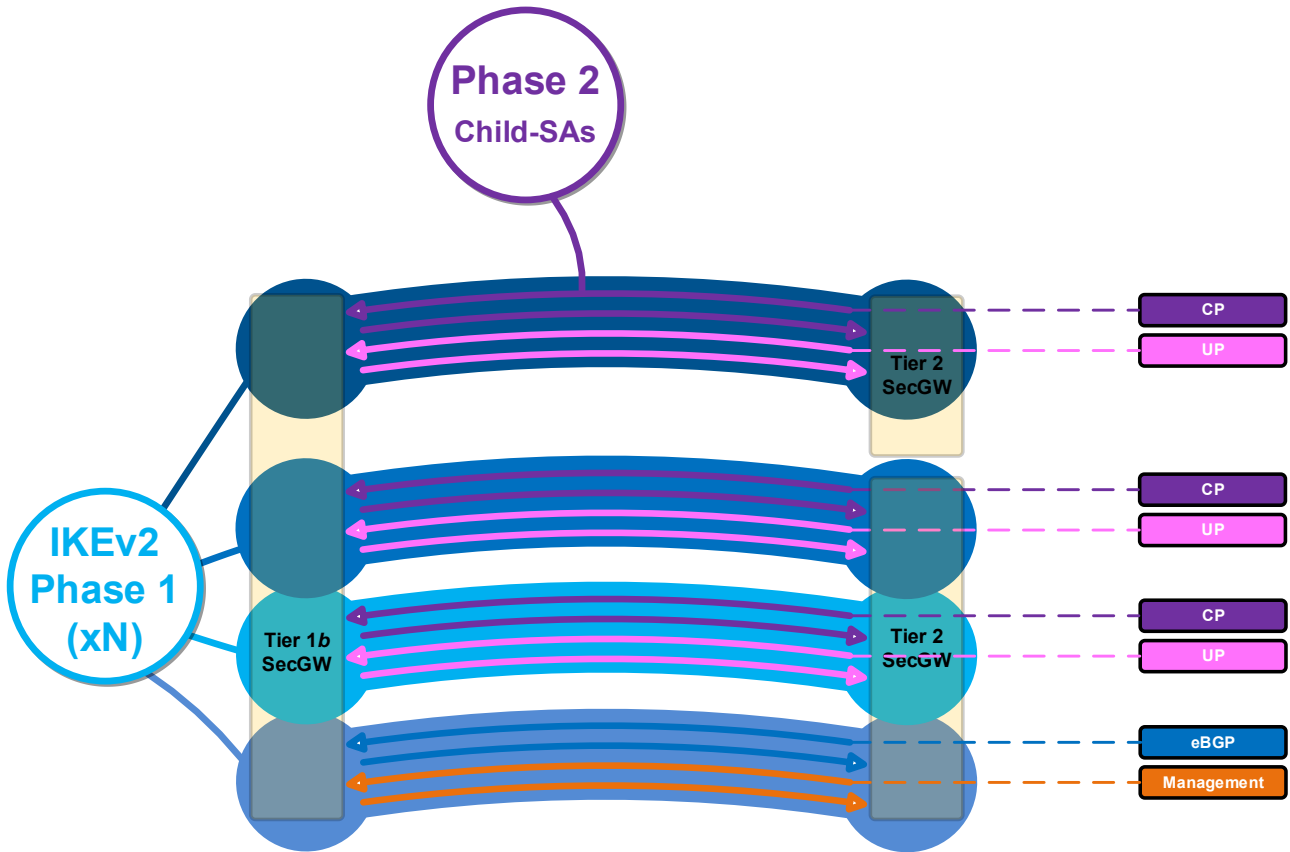


Figure 7-5 – Multiple IPsec Phase 1 – Control/Management Plane and S1/NG Separation across multiple (x N) tunnels.

7.3 Security Controls

- 164. I Security controls are put in place to protect the **Operator Domain** from the **Neutral Host Domain** and equally to protect the **Neutral Host Domain** from the **Operator Domain**.
- 165. M The *Neutral Host* provider must have *Operator* approved physical and logical access controls in place to govern access to equipment and configurations, along with strong audit and Security Information and Event Management (SIEM) processes in place to properly manage in-life configurations (including policy).
- 166. M The *Neutral Host* provider must have strong business processes in place in line with ISO27001 standards to support accuracy of repeatability of all component operational activities.

167. I An overview of the security controls are illustrated in *Figure 7-6*.

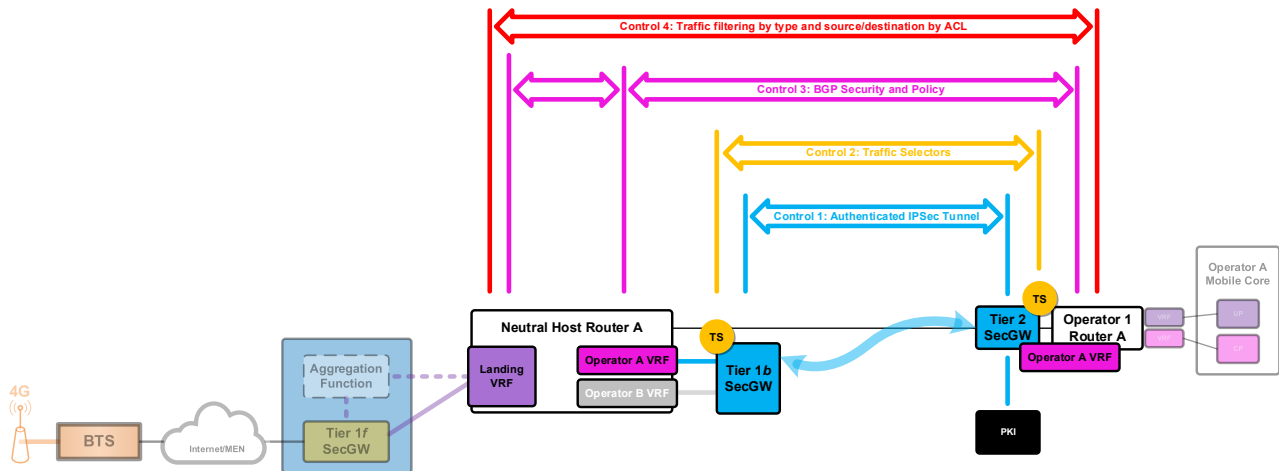


Figure 7-6 – b-interface Security Controls.

168. M The *b*-interface must be secured using the following security controls:

- **Control 1:** Certificate authenticated IPsec Tunnel for *b*-interface;
- **Control 2:** *Operator*-specific Child-SA Traffic Selectors, allowing only expected traffic types and source/destination pairs to traverse the established IPsec tunnel;
- **Control 3:** BGP security and policy to govern exchanged prefixes between the **Neutral Host Domain** and the **Operator Domain** and to control leaking of prefixes to a Landing VRF³² such that **Operator Domain** to **Operator Domain** communication is not routable;
- **Control 4:** Traffic filtering based on traffic types and source/destination pairs as a second line of checks in case of Traffic Selector misconfiguration.

169. I An *Operator* can optionally add a further security control (**Control 5**) in order to provide both SCTP and GTP-U inspection of traffic entering the **Operator Domain** via the aggregated *b*-interface tunnel. This component requires additional equipment and/or licensing to perform this inspection function. Application of SCTP/GTP-U inspection and/or other firewalling applied post Tier-2 SecGW is entirely valid in all cases and is a *per-Operator* decision.

³² Landing VRF component only relevant for the interim Option 2 routing model.

170. I The inclusion of SCTP/GTP-U inspection (Control 5) is illustrated in *Figure 7-7*.

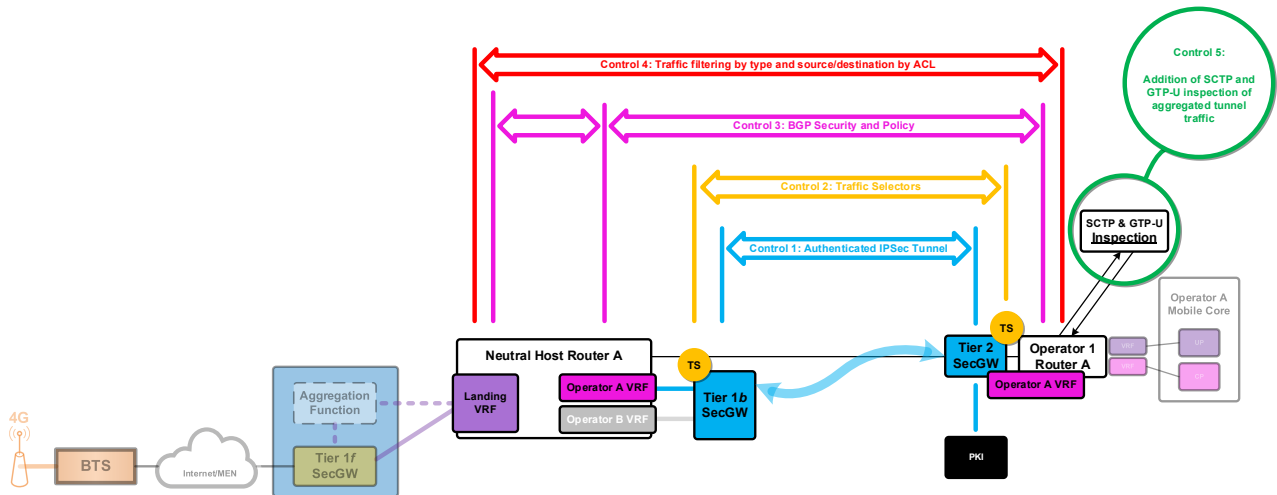


Figure 7-7 – SCTP/GTP-U Inspection.

7.4 Neutral Host PKI

171. M In order to authenticate/authorise connectivity from BTS and Controller nodes, the **Neutral Host Domain** must provide appropriate *Neutral Host* PKI infrastructure.

172. R The certificate management methods that are permitted for the *Neutral Host* PKI include:

- Certificate Management Protocol version 2 (CMPv2);
- Simple Certificate Enrolment Protocol (SCEP).

173. I It is expected that the use of Certificate Management Protocol version 2 (CMPv2) will be *mandated* for *Neutral Host* PKI in a future version of the specification, since it is *strongly* preferred due to a more complete capability of functions such as certificate revocation when compared to other methods such as Simple Certificate Enrolment Protocol (SCEP). However, for initial deployments, the current use of SCEP by some vendor solutions dictates it may be necessary to support both PKI protocols within the **Neutral Host Domain** in order to avoid ruling out deployment of specific vendor solutions.

174. M Regardless of the certificate management protocol used, the following key requirements must be supported:

- Mutual certificate authentication must be used;
- Each BTS and/or Controller entity must hold a unique device certificate;
- The BTS and/or Controller device private key must be generated on the device at the factory;
- The private key must not leave the device;
- The principles of 3GPP TS 33.320 must be followed.

175. R Certificate lifecycle management must be achieved (and proven) and where the capability exists must include Certificate Revocation List (CRL) capability, and ideally Online Certificate Status Protocol (OCSP) capability.

176. R The *Neutral Host* is required to demonstrate the following key components of certificate lifecycle management:

- Enrolment;
- Renewal;
- Replacement;
- Revocation.

7.5 Operator PKI

177. M In order to authenticate/authorise connectivity between the **Neutral Host Domain** and the **Operator Domain** both the *Neutral Host* and *Operator* must provide appropriate *Operator* PKI infrastructure.

178. R The certificate management protocols that are permitted for the *Operator* PKI include:

- Certificate Management Protocol version 2 (CMPv2);
- Simple Certificate Enrolment Protocol (SCEP);
- Certificate Management over Cryptographic Message Syntax (CMC);
- Enrolment over Secure Transport (EST).

Each *Operator* will separately stipulate which of the above they will use.

179. I It is expected that the use of Certificate Management Protocol version 2 (CMPv2) will be *mandated* for *Operator* PKI in a future version of the specification, since it is *strongly* preferred due to a more complete capability of functions such as certificate revocation when compared to other methods. It is noted that the *Operator* PKI may not initially support CMPv2 if the *Operator* has chosen to use SCEP, CMC or EST.

180. M The following must apply to the *Operator* PKI approach:

- Mutual certificate authentication must be used;
- The **Operator Domain** PKI will issue the **Neutral Host Domain** certificate;
- Certificate lifecycle management must be achieved (and proven) and where the capability exists must include Certificate Revocation List (CRL) capability, and ideally Online Certificate Status Protocol (OCSP) capability.

7.6 General Security Considerations

181. M In accordance with RFC2385, the TCP MD5 signature option for cryptographic authentication of both interior and exterior BGP sessions is mandated. The TCP MD5 signature option defines a TCP option for carrying an MD5 digest in a TCP segment, acting like a signature for that segment. Since BGP uses TCP as its transport, it is inherently secure if this mechanism is adopted.
182. R It is recommended that TCP MD5 keys for interior BGP sessions (used internally within the network) should be different to those used for external peering.
183. R HMAC-MD5 cryptographic authentication of IGP and LSP authentication in accordance with RFC3567 (if IS-IS) and ISO 10589 is strongly recommended within both the **Neutral Host Domain** and **Operator Domain**.
184. I The mandating of MD5 will help prevent against risks due to pre-build, non-turn-down, route manipulations and misconfiguration.
185. M MD5 passwords must be obfuscated in configuration views. MD5 passwords will need to be entered into configurations as an unencrypted ASCII key, but in such a way that the unencrypted password cannot be seen within the node configurations.
186. R Per-peer-queuing should be enabled such that separate hardware-based queues are allocated on a per-eBGP-peer basis, such that fair access to shared resources can be granted across all configured BGP peers and to limit the potential impact of attack from/via a specific peer.
187. R Although distant devices spoofing BGP packets is extremely unlikely, the use of TTL-security (in line with RFC 5082) for eBGP peering sessions will provide protection from such attack methods. As the number of hops is known in the environment (i.e. the span between the *Neutral Host* BGP speaker and *Operator* BGP speaker is known), definition of an appropriate TTL-security configuration is easily achieved. Consideration must be given to the multi-hop nature of the eBGP peering connections used for *b*-interface when configuring the TTL-security parameters.
188. R Equipment within the **Retailer Domain** and **Neutral Host Domain** shall be deployed in properly managed sites and racks, with physical security catered for by means of appropriate processes. Access to these locations is to be governed by and is the responsibility of the *Neutral Host*.

189. R Logged administration access to network nodes in both the **Retailer Domain** and **Neutral Host Domain** should be by way of SSH with 2FA employed in an RBAC (Role Based Access Control) model in order to adequately control access to the equipment and avoid the potential for unauthorised reconfiguration and associated compromise (potentially into the *Operator* networks).
190. R Access to the network nodes is to be restricted to specific user accesses.
191. R OSS and network communication access is to be restricted to specific end systems.
192. R Access restrictions shall be configured within the context of Control Plane Policing (CoPP) (which provides rate-limit protection to route processors within the network nodes).
193. M To reduce the risk of DoS attacks directed at the CPU, CoPP must be enabled on external-facing interfaces to allow queuing and discarding of CPU-bound incoming protocol packets should they exceed defined rates. Protection can be applied at the port, VRF interface or ASIC level³³.
194. R CoPP should also be configured such that action is taken as early as possible (i.e. at the line card Network Processor/ASIC level) to discard all packets received for protocols that are not configured on the interface, such that they are not passed to the CPU.
195. R Alarm configuration should be made to enable alarm event generation on breach of defined CPU limits.
196. R Filters must be configured to allow and control BGP protocol traffic (and BFD where used), initiated in either direction, for the *Neutral Host* to *Operator* interfaces.
197. R uRPF (Unicast Reverse Path Forwarding), operated in *loose* mode, should be enabled at appropriate interfaces to provide a level of security against IP-source-spoofing based attacks or misrouting.
198. M Each BGP peer shall be configured with a prefix-limit.
199. R An alarm should be generated at the point that 90% of the BGP prefix-limit threshold is reached, initially with a 'warn only' action at breach. This action on reaching the threshold provides some protection from rogue/spoofed BGP speakers.
200. I Peering session drop due to a BGP prefix-limit being exceeded can be considered by the *Operator*, but this is per-*Operator* security decision.

³³ The most appropriate configuration, along with definition of appropriate rates, will be determined by the specific implementation.

-
201. M eBGP routing policy shall include BGP Prefix Filtering to enforce at both advertisement (i.e. export) and receipt (i.e. import) that *strictly only* the required prefixes to achieve endpoint to mobile core connectivity (and any associated management connectivity) are allowed.
202. R The scope of cleartext traffic will be minimised as far as possible and where traversing links between nodes should be adequately protected by technical means such as MACSEC (where available and practical to deploy) or practical means such as by use of fibre connections within armoured conduit between racks housing the equipment (which must themselves be properly secured).
203. R Where internet-facing and private interfaces exist on a node (for example *f*-interfaces from venues towards Tier1*f* SecGW), separate dedicated physical ports will be required to avoid saturation attack on public facing interfaces also presenting saturation on co-hosted private interfaces. The traffic path from ingress point towards the tunnel endpoints must be fully considered.
204. R Interfaces carrying traffic within the Neutral Host domain, between the Tier1*f* and Tier1*b* SecGW boundaries, shall either be separated on dedicated physical ports and links per operator, or logically separated on shared physical links, noting that where shared physical links are used, appropriate QoS policy must be in place to ensure fairness of traffic for the multiple MNO services.
205. R Where virtualised elements are used (i.e. VNFs), for interfaces requiring physical port separation must have distinct VNF -> vNIC port -> pNIC port mappings, either via dedicated vSwitches or via technologies such as SR-IOV.
206. M Appropriate capacity management of both external and internal links must be in place in all domains.

207. I An overview of port separation requirements can be seen in *Figure 7-8* below:

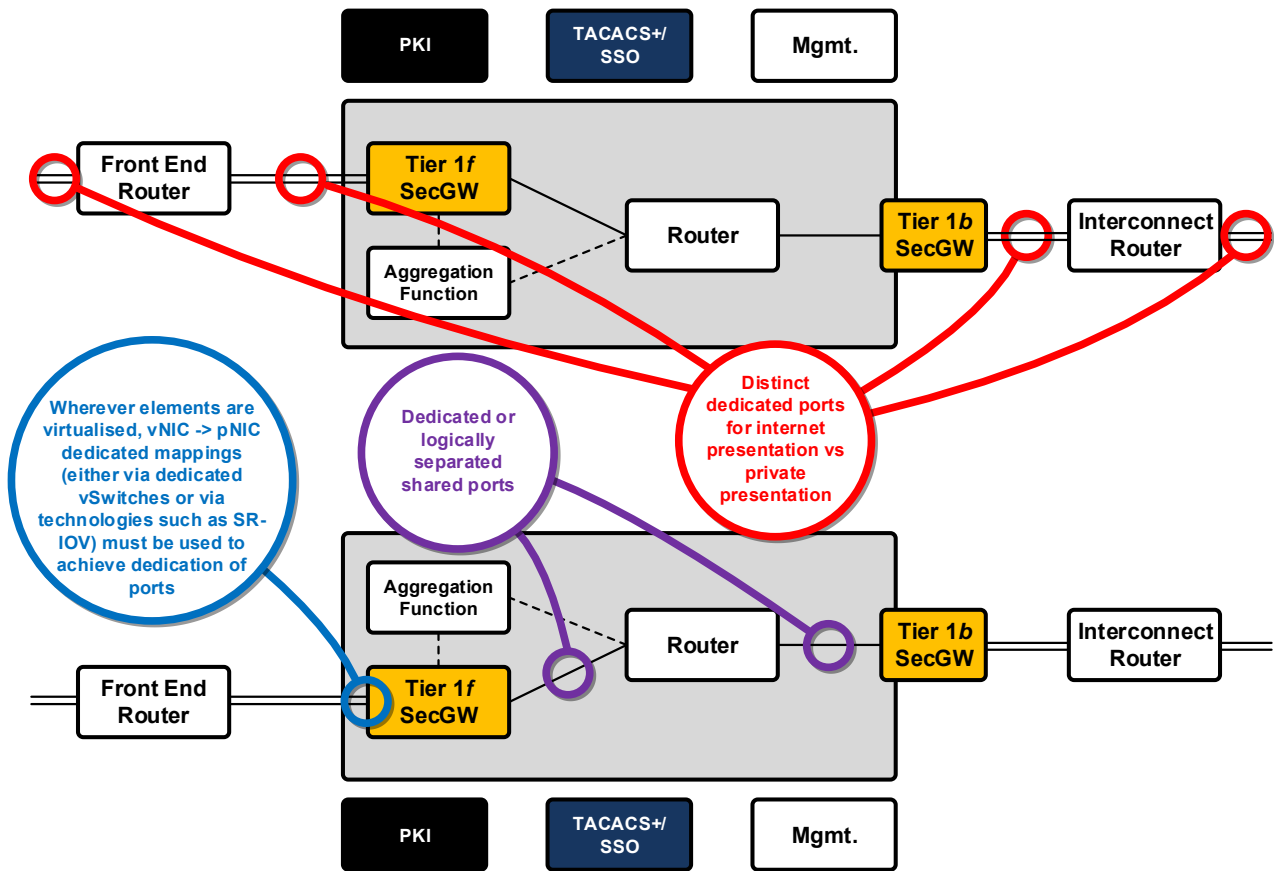


Figure 7-8 – Port Separation Requirements.

7.7 DoS/DDoS Protection

208. I For Internet connectivity scenarios, tunnel endpoint public addressing (whether *f*-interface or *b*-interface) will be (by necessity) routable from the Internet and therefore become part of the DoS/DDoS attack surface.
209. R In order to avoid loss of service (potentially affecting all *Operators*), it is strongly recommended that DoS/DDoS protection against volumetric attacks is in place to protect the (various) exposed public interfaces.
210. I DoS/DDoS protection is the responsibility of the *Neutral Host* for connectivity into the centralised Controller nodes, the Tier-1*f* SecGW and the Tier-1*b* SecGW, whilst being the responsibility of the *Operator* for connectivity into the Tier-2 SecGW.
211. R DoS/DDoS protection should not add significant latency, particularly to the *f*-interface which may be intolerant of increased latency. In the event of a DoS/DDoS protection induced latency increase, the behavior of the BTS solution must be properly understood in terms of customer experience impact.

212. R Since traffic is IPSec encrypted for exposed external interfaces, perimeter-based filters (e.g. router ACLs) to limit traffic to expected IPSec flows (e.g. ESP, ISAKMP, NAT-T UDP 4500) should be applied in each domain.
213. R Since the IPSec encrypted flows cannot be inspected by a mitigation platform, protection should concentrate on volumetric attack protection, unless in-line platforms capable of heuristic based encrypted flow assessment form part/all of the anti-DoS/DDoS solution.
214. I From the *Neutral Host* perspective, the DDoS mitigation approach should be understood in terms of BGP diversion routing (if applied) to support mitigation (e.g. where a /24 range is redirected for DDoS inspection) and as such consideration should be given to the addressing structures per service in order to avoid diversion of un-related (non-NHIB) services.

7.8 GMLC Location and Radio Location Lock

215. M In order to support license obligations relating to emergency calling, it is necessary for appropriate zone code and latitude/longitude data for deployed cells to be available via Gateway Mobile Location Centre (GMLC). As such, development to ensure this can be correctly presented for the NHIB solution at the OSS layer and GMLC will be required per *Operator*.
216. R Consideration must be properly given to the capability of the solution in terms of presenting location data, due to the aggregation approaches used. For example, the hosting of radio nodes at multiple distinct venues via a shared Controller, which must, by definition, be then able to adequately define radio nodes at a logical level that can be mapped to location data for a specific venue, not simply for the Controller as a whole.
217. I Given that the deployment model for BTSs will be via the **Retailer Domain** and **Neutral Host Domain**, the *Operator* will not have direct control over either the initial deployment or subsequent movement of BTS devices. Given that supporting equipment (e.g. PoE switches and CPE) are quite likely to get moved with the BTSs (e.g. if an enterprise moves premises and does a lift and shift of their infrastructure) and given that public IP reachability will still be possible from a new BTS location where *f*-interface is Internet presentation, unauthorised and undetected movement of small cell BTS presents a risk to correct emergency call location (and therefore license obligations).
218. M In order to prevent the movement of small cell BTSs from their authorised location, radio location lock functionality is mandated.
219. R BTS nodes analysing the surrounding macro environment (e.g. MNC Id or network measurements) at the point of deployment should be able to present an appropriate alarm (with configurable

down-action) to alert of a potential equipment move on changes of the in-life visible macro environment.

220. M Where public-IP addressing is used for BTS connectivity, IP based geo-location to ensure that BTS nodes are not taken out of the country or region of initial registered installation is mandated.

8 QUALITY OF SERVICE

- 221. I In order to simplify the reality of interconnecting to several *Operator* networks, it is required that the *Neutral Host* is considered the anchor-point for the QoS scheme.
- 222. M It will therefore be the responsibility of the *Neutral Host* to provide details of their deployed QoS scheme, such that the end-to-end QoS approach can be appropriately defined and implemented per *Operator*.
- 223. I Four QoS domains exist within the NHIB architecture, as illustrated in *Figure 8-1*.

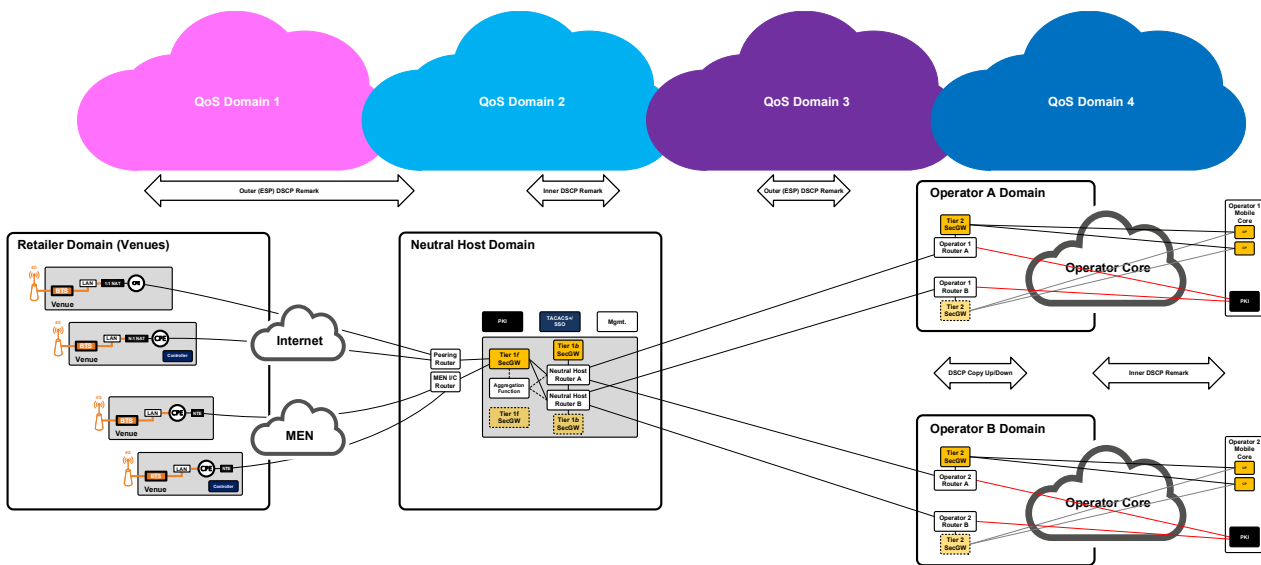


Figure 8-1 – QoS Domains.

- 224. R **QoS Domain 1** - Given that different *Neutral Hosts* will exist and may each use a number of different services to connect the **Retailer Domain** to the **Neutral Host Domain**, where those differing services may themselves have different QoS requirements and constraints, it is the responsibility of the **Neutral Host Domain** and **Retailer Domain** to appropriately classify, treat and remark traffic in order to ensure its fair treatment between different *Operators*.
- 225. R **QoS Domain 2** - Once within the **Neutral Host Domain**, any further classification, treatment and remarking of traffic to appropriately control its behavior within the **Neutral Host Domain** itself is clearly the responsibility of the *Neutral Host*. It should be noted that due to the 'air-gap' introduced within the **Neutral Host Domain**, the opportunity to remark the *inner* DSCP rather than the *ESP outer* exists in this domain and may be required in order to achieve alignment with the differing QoS schemes deployed in the **Operator Domains**.
- 226. R **QoS Domain 3** - Transit services between the **Neutral Host Domain** and the **Operator Domain** should ideally be dark-fibre or wavelength services, without any specific QoS constraints, but where an alternative service is used, the requirements of that service must be met and may require specific classification, treatment and remarking of traffic. This must be achieved at both the sending and

receiving sides of this QoS domain (from the perspective of both directions), to align with the requirement of the transit service and also to realign from that back to the required *received* traffic QoS model. This is therefore a shared responsibility between the **Neutral Host Domain** and the **Operator Domain**.

227. R **QoS Domain 4** - Once within the **Operator Domain**, appropriate remarking and classification will be required in order to align with the *Operator* core (and mobile core) QoS schemes.
228. R Given the reality that much of the remarking possible is only possible at the ESP *outer* packet level due to the IPSec nature of the flows, and given that features such as IPSec 'copy-down' in the decryption path (as well as the standard 'copy-up' in encryption path) may not be universally available, the actual end-to-end QoS approach will need varying approaches in different scenarios.
229. M The detail of QoS marking and remarking must be the subject of adequate low-level design consideration at the point of implementation for each *Neutral Host*, with considerations made in all the above QoS domains.
230. I Where possible, the amount of specific treatment within the **Neutral Host Domain** will be minimised and standardised, to reduce complexity within that domain, but the nature of the QoS requirements means that this will not always be possible. As such, any solution must contain the capabilities necessary to cover the QoS requirements stated for all QoS Domains.
231. M The **Retailer, Neutral Host** and **Operator Domain** components of the solution must together contain the capabilities necessary to cover the QoS requirements stated above.
232. R H-QoS capabilities are to be deployed where necessary to achieve fairness in treatment of equivalent traffic classes between *Operators*. It should be noted however that appropriate dimensioning of the solution may dictate H-QoS unnecessary – although where this is believed to be the case, H-QoS should be retained as a remedy measure.
233. R It is required that some form of Call Admission Control is applied or that a planning and monitoring approach is used to avoid the scenario where excess traffic (beyond the CIR=PIR provision for the EF queue) on a *f*-interface link leads to a degradation for calls in progress.
234. R Appropriate 802.1p marking is required to be applied at the hosted BTS, along with DSCP marking, in order allow prioritised traffic handling within the switch infrastructure at the venue.

9 End-to-End Management

235. I An overview of the high-level requirements for multi-tenancy management is illustrated in *Figure 9-1*.

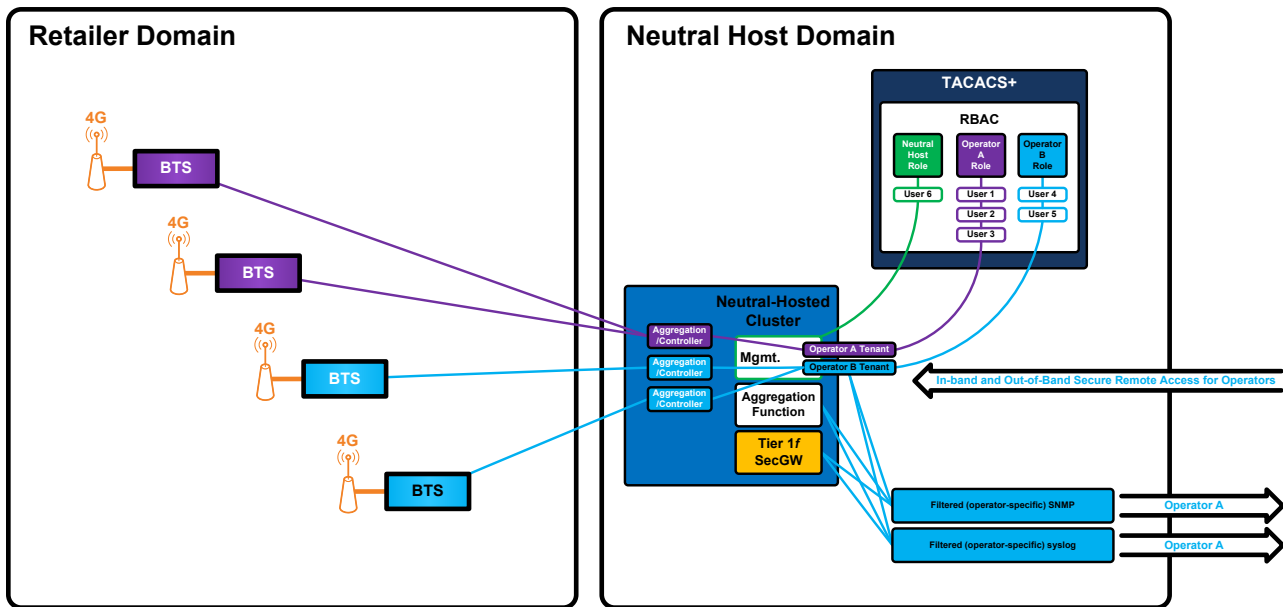


Figure 9-1 – Management Multi-Tenancy.

236. R Since the *Neutral Host* will provide management capability for the **Neutral Host Domain** and **Retailer Domain** deployed BTS (and related) equipment, with these providing service to the **Operator Domain**, it is required that *Operators* have some agreed level of visibility and management capability within the **Neutral Host Domain**.
237. R Given that configuration data and event/alarm data (i.e. syslog and SNMP) will be present on a shared platform but relating to multiple *Operator* deployments, the management capability will be required to support multi-tenancy.
238. R In order to drive this multi-tenancy (and to meet TSCoP requirements on AAA and management separation), it will be necessary to control access to tenants within the multi-tenant management by means of RBAC (Role Based Access Control) with 2FA (Two Factor Authentication).
239. R It is required that SNMP alarming and syslog event visibility must be filtered such that only those events relevant to each tenant are visible from within that tenancy.
240. M In-band management must be provided in line with the routing, resilience and IPSec models detailed in Section 4.3 and Section 7.2.

241. M Out-of-Band management must be provided, with this typically expected to be carried via IPSec over Internet transit to an OOB-specific tunnel endpoint within the **Neutral Host Domain**.
242. M Both In-Band and Out-of-Band management methods must be subject to role-based access control (RBAC) and two-factor authorization (2FA).

--- End of Document ---