Broadband Forum Liaison To:
IETF Service Function Chaining (SFG) WG
Jim Guichard, jguichar@cisco.com
Thomas Narten, narten@us.ibm.com
IETF Liaison Statements, statements@ietf.org

From:
Christophe Alter, Broadband Forum Technical Committee Chair (christophe.alter@orange.com)
Sultan Dawood, Broadband Forum Marketing Committee Chair (sudawood@cisco.com)

Communicated by:
David Sinicrope, Broadband Forum Liaison Officer to IETF (david.sinicrope@ericsson.com)

Date: 17 February 2014

Subject: Broadband Forum Work on Flexible Service Chaining (SD-326)

Dear Jim and Thomas:

We were informed by the Broadband Forum (BBF) Liaison Officer to IETF (David Sinicrope) that the IETF had recently formed the Service Function Chaining WG. In reviewing the WG charter, we thought there might be IETF interest in our active project on Flexible Service Chaining. We would like to share information regarding the scope of this BBF project and explore any potential coordination with the IETF.

First as background, the BBF Flexible Service Chaining work is being done in our Service Innovation and Market Requirements (SIMR) Working Group (WG). The mission of SIMR WG is to drive BBF with medium-to-long term work directions and requirements in order to lead innovation in broadband networks. For additional background on the BBF and on the different BBF Working Groups missions, please visit the following: http://www.broadband-forum.org

The current Flexible Service Chaining project is targeted to deliver a Study Document (SD-326), with the final document intended for internal BBF use; the goal of SD-326 is to select use cases and identify technical gaps that will require technical specification work in other BBF WGs (e.g., End-to-End Architecture WG) or partner organisations (e.g. IETF). We expect that it will lead to published extensions to broadband network architecture and corresponding nodal requirements (a typical BBF activity) and in turn, possibly imply some new protocol(s) or extensions (a typical IETF activity). SD-326 is a work in progress and is expected to complete in the 2014 timeframe.

Some contents from SD-326 (e.g., purpose, scope, representative use cases) are copied below. Although SD-326 is intended for internal BBF use, we felt it may be of interest to the IETF to help inform the IETF of related BBF work on Flexible Service Chaining and help facilitate positioning and any potential coordination of work between our two organizations.

At your upcoming IETF SFC WG meeting we expect that Hongyu Li (SD-326 Co-Editor) will be participating and can present this liaison statement in addition to answering questions about the work.

The BBF welcomes input and coordination on this project from the IETF and looks forward to hearing more about the work of the SFC WG. We will continue to keep you informed of progress on our Flexible Service Chaining project. For information, the upcoming quarterly Broadband Forum meetings are listed at the end of this liaison.

Sincerely,

Christophe Alter,
Broadband Forum Technical Committee Chair

Sultan Dawood,
Broadband Forum Marketing Committee Chair

CC:
Christophe Alter, BBF Technical Committee Chair (christophe.alter@orange.com)
Sultan Dawood, BBF Marketing Committee Chair (sudawood@cisco.com)
George Dobrowski, BBF SIMR WG Co-Chair, georgedobrowski@mail01.huawei.com
Michael Fargano, BBF SIMR WG Co-Chair, michael.fargano@centurylink.com
Christele Bouchat, BBF SIMR Vice-Chair, Christele.bouchat@alcatel-lucent.com
Robin Mersh, BBF CEO, rmersh@broadband-forum.org
Hongyu Li, BBF SIMR WG SD-326 Co-Editor, hongyu.li@huawei.com
Jerome Moisand, BBF SIMR WG SD-326 Co-Editor, jmoisand@juniper.net
Gabrielle Bingham, BBF Secretariat, gbingham@broadband-forum.org
Adrian Farrel, IETF Routing Area Director, adrian@olddog.co.uk

Date of Upcoming Broadband Forum Meetings

A detailed list of upcoming meetings can be found at http://www.broadband-forum.org/meetings/upcomingmeetingsataglance.php

Purpose and Scope (from DRAFT SD-326_Rev03):

Purpose

In order to support business and residential, fixed and mobile, wholesale and retail markets, TR-144 described various requirements including the need for network interconnection standards for broadband access, QoS support, Bandwidth on demand, increased overall bandwidth, higher network reliability and availability.

New ways of defining services are required to keep up with market needs, seeking more flexibility in service deployment, faster service feature delivery, increased automation, elastic service bursting, etc.

Service chaining allows complex services to be created out of simpler service-enabling elements through composition, e.g. stringing service points together while possibly constraining the corresponding data path.

The output of this project will provide guidance to BBF's Technical and Marketing Committees on the work needed to bring flexible service chaining concepts to the level of detail necessary to define broadband network element requirements for implementation. This Study Document will also provide a reference for other service chaining standards organizations.

Scope

WT-178 issue 1 built on TR-144 and TR-145 to establish the basic architectural principles and protocols to implement a multi-edge architecture and Layer 2 and 3 session control, including the hierarchical-BNG construct, stringing together a Broadband Network Gateway and a Broadband Service Gateway by L2 forwarding means (e.g. MPLS).

WT-178 issue 2 will augment WT-178 issue 1 by specifying requirements for BNGs and BSGs to address requirements for basic service chaining. Service chaining being the concept of stringing Service Enforcement Points (aka middlebox) together that deliver a set of services. Basic service chaining will be limited to "pre-compiled" and "pre-provisioned" chains that transit a single BSG in addition to the (edge) BNG.

This Study Document intends to study market requirements and use cases for Flexible Service Chaining, as well as suggestions for new project after analysis of gaps with existing projects in BBF.

This project will go for more flexible forms of service chaining going beyond the simple goals covered by WT-178 Issue 2, and will investigate some or all of the following topics:

- Multiple forms of service enforcement at L3 or above (e.g. DPI, Firewall/Security, Parental Control, Captive Portals, etc), on various types of network systems (wireline/mobile/WiFi gateways, service routers, dedicated appliances, virtual machines, etc);
    - Identify a collection of business-enabling use cases
- Complex service chains, e.g. more than two Service Enforcement Points:
    - Service chains with a fixed or variable shape (e.g. open or closed chains; dynamic provisioning of service chains; dynamic changes to the data path for a given traffic session/flow; symmetric or asymmetric forwarding);
    - Intra-facility or cross-facility (e.g. Central Office to Point of Presence to Data Center) service chains;
    - Intra-domain or cross-domain (e.g. multiple service providers involved, or loosely coupled groups within a given service providers) service chains;
    - Address interface mobility when service enforcement is supported by a virtual machine, in order to support virtual machine migration
- Carrier class service chains:
    - Distributed QoS enforcement along the service chain, per session and per subscriber;
    - Define network service SLAs in terms of bandwidth, latency, OAM for service chaining.
- Provide high level of automation for intra and cross-domain scenarios:
    - Fully subscriber-aware session authorization & accounting that may or may not necessarily require stateful and tightly coupled L3 Session Control
    - Dynamic forms of load-balancing and high-availability, e.g. based on state information being exchanged
    - Possible support for elastic bandwidth on demand.

Example Use Cases and Input to Market Requirements (from DRAFT SD-326_Rev03):

1. Use Case Template

Title	*Short title, reminiscent of key aspects of the service being provided and the corresponding service chain.*
Service Model and Story Highlights	*Service model: short description of the service(s) as perceived by end users. What type of end users are targeted (e.g. consumers, businesses, other service providers).* *Story highlights: a few key points to further characterize the service and the way it should be provided.*
Business Drivers	*Description of business drivers, typically using two perspectives:* *1. external-facing: business considerations towards end users (e.g. sell a new*

	service) *2. internal-facing: business considerations internal to the SP (e.g. supplier management considerations, Capex/Opex considerations)*
Deployment Model	**The service chain(s) may include *XX* primary Service Functions on the data path:** *Brief description of set(s) of service functions to be chained.* **Geographical distribution**: *brief description of physical facilities involved (e.g. hosting service functions) and how geographically distributed they are.* **Administrative boundaries**: *brief description of administrative domains involved (e.g. service providers and/or independent organizational groups within) and who does what.*
Actors	*A list of network(s) involved, service providers involved and types of end users using the services being described. Something like:* Multiservice Broadband Network Broadband Service Provider Broadband User (consumers)
High-level architectural context	*(Instruction to contributors): start by a simple and easy to understand drawing (please try to stay consistent with drawings from existing use cases), WITHOUT implying any detailed architectural choice.* *More detailed description of expected behavior of each service function, and the way they chain/interact.* Architectural attributes (functional view): *(Instruction to contributors): although optional, please try to stay in the structure of these attributes listed below. If some key information doesn't fit, please add to the end of the list.* **Service chain shape**: *chain of X systems, shape of the chain (open/closed, line/graph, static/dynamic, etc). Service chains will typically start by a BNG (primary BBF service function in TR-101), but does not have to.* **Performance**: *SLA considerations, intra-network performance considerations* **Load balancing & Resiliency**: *more or less automated, static or dynamic forms of load-balancing (+ criteria). Redundancy considerations (e.g. primary/backup, 1:1 vs N:1 model, geo-redundancy, etc).* **Automation & Lifecycle:** *emphasis suggested on OSS integration and service provisioning & activation, and ways to automate & reduce related costs and processes. Two levels to elaborate on: the management & lifecycle of service chains, and the management & lifecycle of service instances (e.g. user sessions).* **Traffic engineering**: *network traffic engineering considerations (e.g. service tunnels and related network paths). Static vs dynamic.*

	Symmetrical forwarding: *required, not required? why?* **Other considerations (e.g. mobility, etc)**: *as needs be, open-ended.*
Related and Derivative Use Cases	**Existing use cases**: *possible relation to existing SD-326 use cases*. **Derivative use cases**: the use case described in this table might morph into something slightly different over time, i.e.: • *New service functions, new shape of the service chain* • *Different geographical/administrative distribution, etc.*
Issue(s) Spotlight	*(Instruction to contributors): if the use case presents a challenge (or more) that could benefit from more explanations, please use such an entry to elaborate, optionally including supporting pictures.*
Inputs to market requirements	*(Instruction to contributors): this entry allows to summarize in an informal manner (e.g. bulleted list) key points which are suggested to be translated in market requirements in the corresponding section of SD-326.*

2. Internet access, CGNAT and Web Filtering

Title	Internet access, CGNAT and Web Filtering
Service Model and Story Highlights	Service model: the Service perceived by Broadband Users (consumers) is Internet access, with optional Web URL filtering (e.g. parental control) and a single private IPv4 address being allocated. Story highlights: For service users selecting a Web filtering "null" profile, and for non-HTTP traffic in any case (notably latency-sensitive traffic like VoIP), the Broadband Service Provider wishes to improve performance to bypass the HTTP Filtering function for corresponding traffic. To address the NAT requirement (which stems from a planned drought of public IPv4 addresses, and challenges in making an IPv6 migration timely happen), the SP wishes to avoid any change to the consumer network, therefore to perform IP address translation (e.g. port-based) inside the MBN.
Business Drivers	The business value of the Service is to satisfy a class of concerned customers who may want to prevent their family from accessing 'inappropriate' Web sites (different profiles of filters being available, including no filtering at all), while providing a high performance traditional IPv4 service in a time of IP address space crunch. The URL filtering must not impact the latency-sensitive traffic like VoIP. In addition, there is a desire to use best-of-breed devices from separate vendors hosting Service Functions, e.g. a traditional BNG (e.g. router-based), combined

	with a flexible HTTP Filtering system (e.g. server-based), combined with a CGNAT function system (e.g. router-based or server-based).
Deployment Model	**The service chain(s) may include three primary Service Functions on the data path:** 1. Regular BNG (L3) function 2. HTTP URL Filtering function (optional) 3. Carrier-grade NAT function For non-HTTP traffic or for service profiles without Web filtering, a simpler service chain is used, including two primary Service Functions on the data path: 1. Regular BNG (L3) function 2. Carrier-grade NAT function **Geographical distribution**: it is assumed that physical systems hosting all 3 service functions are collocated in a single facility (e.g. in a Metropolitan Point-Of-Presence). **Administrative boundaries**: all Service Functions are contained within a single MBN "+" network, operated by a single service provider and administrative entity.
Actors	Multiservice Broadband Network Broadband Service Provider Broadband User (consumers)
High-level architectural context	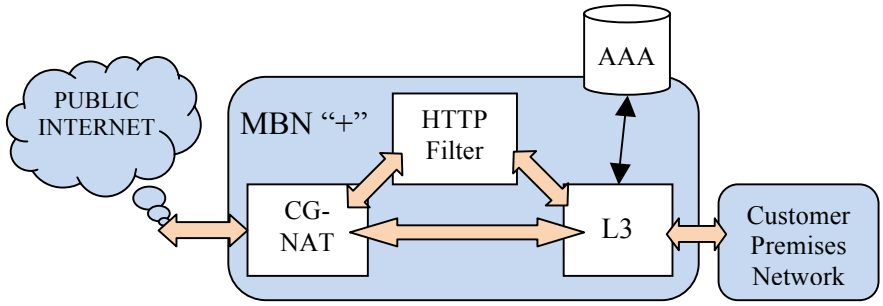 All Service Functions are deployed as part of an extended Multiservice Broadband Network ("MBN +"), allowing user traffic to flow to/from the public Internet, while being appropriately serviced. The BNG is configured to allocate private IPv4 addresses to Residential Gateways. Two types of service chains are used, either BNG to CGNAT, or BNG-to-HTTP-Filter-to-CGNAT. User traffic is selectively routed to the appropriate service chain, depending on the service terms (e.g. as provided by a AAA server) and the nature of the traffic (e.g. HTTP or not). Depending on the exact terms of the service being subscribed to, the BNG may enforce various types of QoS/Policy profiles, and the HTTP Filter may enforce various types of filtering

	list. The CGNAT function typically enforces the same processing to all traffic, with no service-based differentiation. Architectural attributes (functional view): **Service chain shape**: chain of 2 or 3 systems; straight (open) line **Performance**: low-latency traffic (e.g. VoIP) not impacted by HTTP Filtering; SLAs similar to typical Internet access consumer services (e.g. tiered peak rate, no committed rate). **Load balancing & Resiliency**: automated balancing between a farm of servers supporting the Web filtering function for a given BNG, based on simple criteria (e.g. max number of user sessions per server). Primary/backup (1:1 redundancy) system hosting the CGNAT function. Traffic associated with a given user session must always traverse the same instance of Web filtering and CGNAT Service Functions. **Automation & Lifecycle:** static configuration of service chains through regular OSS systems, no need for further flexibility required in this case. Choice of exact service profile per user fully driven by an external AAA system, with minimum OSS integration burden and changes compared to a regular BNG-only service deployment, enabling full automation of service instantiation per user session. **Traffic engineering**: statically configured bandwidth provisioning and traffic engineering. **Symmetrical forwarding**: required for all traffic **Other considerations (e.g. mobility, etc)**: none
Related and Derivative Use Cases	**Existing use cases**: no existing SD-326 use case appears to be directly related to this one (so far). **Derivative use cases**: the use case described in this table might morph into something slightly different over time, i.e.: • The Web Filtering and CGNAT functions might be located in a more centralized data center, • There might be a need for a deeper inspection Service Function for some forms of HTTP content, • IPv6 traffic may require Web Filtering, but no CGNAT, • Etc
Issue(s) Spotlight	The fact that the service chain is terminated by a CGNAT function implies that traffic using private IP addresses need to be conveyed across the service chains. As such address space is limited, it makes desirable to use IP VPNs as a way to logically segment traffic and to allow overlapping IP address spaces. This implies in turn that service chains (and Service Functions) need to support VPN awareness. On the other hand, the use of CGNAT greatly facilitates the symmetry of

	upstream/downstream traffic through the service chain, as routing protocols on the core network will naturally direct downstream (network to user) traffic associated with public IP address subnets to the CGNAT Service Function, where the traffic can then be steered along the service chain towards the end user premises.
Inputs to market requirements	This use case illustrates several key needs for service chaining: - The first Service Function (BNG) needs to be capable of classifying traffic (including for a given user session) before steering to the appropriate service chain - The association between traffic associated with a given user session and a given service profile needs to be somehow propagated between the first Service Function (BNG) and the other Service Functions - Load balancing and steering of traffic for a given user session on a given service chain needs to be deterministic enough, so that all traffic for a given user session always traverse the same instance of Service Functions. - VPN awareness is required as part of the service chaining infrastructure.

3. Internet access and network-based Parental Control service

Title	Internet access, Parental Control, No CGN, IPv6
Service Model and Story Highlights	The Service Model as perceived by Broadband Users (consumers) is Internet access with an optional network-based Parental Control service, in a context where the residential gateways are assigned a public IPv4 address and possibly a global IPv6 prefix. For service users selecting a Parental Control "null" profile, and for non-HTTP traffic in any case (notably latency-sensitive traffic like VoIP), the Broadband Service Provider wishes to improve performance to bypass the Parental Control platform for corresponding traffic. The Parental Control function not only inspects the incoming URL requests from the Broadband Users (upstream) but also performs a deep analysis on the contents returned by the web servers (downstream), in order to, for example, block/filter/alter some undesirable objects contained in the resulting web page, such as inappropriate pop-ups, banners, etc. In this use case, the SP is not pressured by the public IPv4 address exhaustion and thus does not need to introduce a NAT function in its network. All Broadband User gateways are provided with a public IPv4 address and a global IPv6 prefix, regardless of whether they have subscribed to the Parental Control service or not.
Business	The business value of the Service is to satisfy a class of concerned customers who

Drivers	may want to prevent their family from accessing various categories of Web sites from any device in the LAN (URL filtering), as well as from unexpectedly encountering undesirable contents (discussions on social networks, Ads pop-ups, unexpected URL redirections (deeper inspection). There is a desire to use best-of-breed service-enabling devices from separate vendors, e.g. a traditional BNG (e.g. router-based), combined with a flexible Parental Control system (e.g. server-based).
High-level architectural context	**The service chain(s) may include two primary service-enabling functions on the data path:** 1. Regular BNG (L3) function 2. Parental Control function For non-HTTP traffic or for service profiles without Parental Control or for walled garden HTTP traffic (e.g.: NSP's VoD, NSP's music streaming, etc.), the traffic is simply forwarded normally by the BNG. **Geographical distribution**: the BNG and the Parental Control function may in some cases be collocated in the same premise, but are likely to located in different premises in most deployments: Parental Control platforms are designed to support a large number of subscribers compared to the number of customers connected to a given BNG who would have subscribed to the parental control service. **Administrative boundaries**: all service-enabling functions are contained within a single MBN "+" network, operated by a single service provider and administrative entity.
Actors	Multiservice Broadband Network Broadband Service Provider Broadband User (consumers)
High-level architectural context	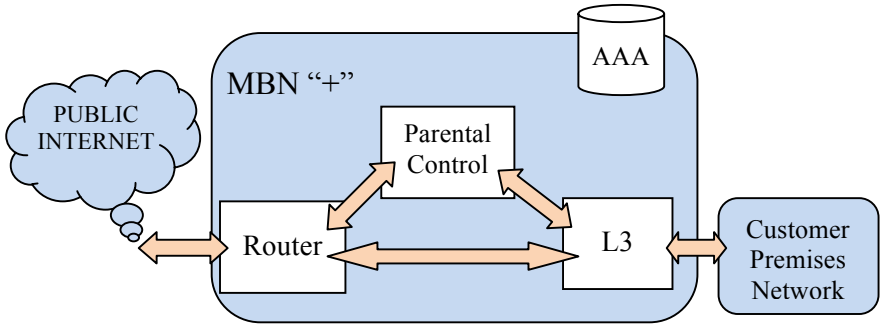 All service-enabling functions are deployed as part of an extended Multiservice Broadband Network ("MBN +"), allowing user traffic to flow to/from the public Internet, while being appropriately serviced. The BNG is configured to allocate public IPv4 addresses to Residential Gateways

	as well as global IPv6 prefixes. Two types of service chains are used, either BNG to Internet, or BNG-to-Parental-Control-to-Internet. User traffic is selectively forwarded according to the appropriate service chain, depending on the service terms (e.g., as provided by a AAA server to the BNG) and the nature of the traffic (e.g., HTTP or else). Depending on the exact terms of the service being subscribed to, the BNG may enforce various types of QoS/Policy profiles, and the Parental Control may enforce various types of filtering list. Architectural attributes (functional view): **Service chain shape**: only 2 functions – the Parental Control Platform and the BNG. **Performance**: low-latency traffic (e.g. walled-garden VoIP) not impacted by the Parental Control function; SLAs similar to typical Internet access consumer services (e.g. tiered peak rate, no committed rate). **Load balancing & Resiliency**: automated balancing between a farm of servers supporting the Parental Control function for a given BNG, based on simple criteria (e.g. max number of user sessions per server). **Automation:** choice of exact service profile per user fully driven by an external AAA system, with minimum OSS integration burden and changes compared to a regular BNG-only service deployment. **Traffic engineering**: statically configured bandwidth provisioning and traffic engineering. **Symmetrical forwarding**: required for all traffic especially in the case of this Parental Control function which also inspects the content of the Web downstream traffic. **Filter granularity & lifecycle**: at the BNG level, traffic per user session, and HTTP vs non-HTTP must be appropriately filtered and classified for steering on the proper service chain. Such filters will have the same lifecycle as a user session (e.g. PPP, DHCP). **Other considerations (e.g. mobility, etc)**: none
Related and Derivative Use Cases	**Existing use cases**: no existing SD-326 use case appears to be directly related to this one (so far). Derivative use cases: None.
Issue(s) Spotlight	**Issue 1: Symmetrical Traffic** In the absence of a NAT function for IPv4 traffic, how can the SP ensure that the downstream web traffic is going to be forwarded to and processed by the Parental Control function? How can the SP ensure that the downstream web IPv6 traffic is going to be forwarded to and processed by the Parental Control function?

Description of the issue:

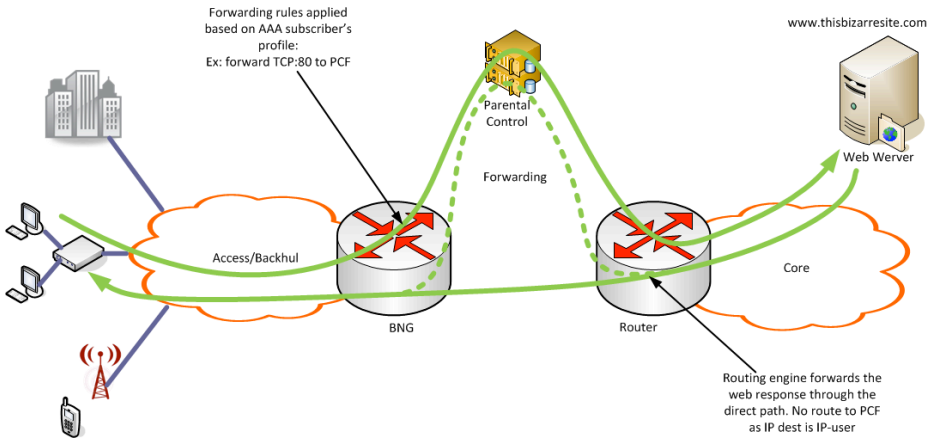

A web request is issued by a Broadband User who has subscribed to the Parental Control service to the web site www.thisbizarresite.com. An IPv4 datagram with `IP-source = IP-user` (after NAT on RG) and `IP-dest = IP-www.thisbizarresite.com` is intercepted by the BNG which forwards it to the PCF. The PCF inspects the URL, and in this case normally forwards the datagram. The Web server replies with a set of datagrams with `IP-source= IP-www.thisbizarresite.com` and `IP-dest = IP-user`. The path to reach the user is the direct path. Hence the response is not seen by the PCF and thus can not be inspected.

Issue 2: Multiple filtering profiles and Subscriber's Parental Control filtering specific configuration

The NSP may propose several generic parental control profiles to its subscribers, to filter identified categories of contents (politics, drugs, army, pornographic, etc). In addition, the NSP may propose to the subscriber to "tune" its filtering profile (adding/removing URL for example).

One issue to address is how the Parental Control selects the filter and possibly the specific configuration of the subscriber.

Purpose of mentioning this issue is to wonder if it would be wise to envision the circulation of some information specific to each subscriber throughout the service chains to pass some information to the functions in the chain if needed.

Issue 3: Encrypted HTTPS traffic

Encrypted web traffic (https) already represents a very significant part of Web traffic and is likely to shortly become the main or even the only method to carry Web data over the Internet.

	In this context, the Parental Control platform must be able to scan inside SSL, to eventually filter some undesired content if any. This issue needs to be addressed for the Parental Control service to operate properly, but solving that issue is not in the scope of SD-326.
Inputs to market requirements	This use case illustrates several key needs for service chaining: • Need to enable symmetrical traffic along a service chain for a given flow, • Need to circulate some subscriber specific information to appropriate service functions.

4. Lawful Intercept and Deep Packet Inspection

Title	Lawful Intercept and Deep Packet Inspection for broadband users
Service Model and Story Highlights	Lawful Intercept is not a service offered to broadband users per se, but it can be considered as a service provided to an external Government Agency and it represents a need from the Network/Service Provider point of view with reference to legal obligations. Deep Packet Inspection is usually used to perform traffic statistic analysis but in a medium and long term future it could be also offered as a service to broadband users (e.g. per subscriber or per application quota based services). The Service Model offered as perceived by broadband users is VoIP, Video and Internet access and connectivity can be L3 VPN or IP. The traffic of the customers for whom it is not necessary to perform statistic analysis through the Deep Packet Inspection service function goes straight and does not traverse the relative service enabling function. The Lawful Intercept service function is traversed only based on the obligation from an external Government Agency to intercept the traffic from and to a specific customer. When the need to intercept traffic from and to a specific customer occurs, at the BNG level this traffic is replicated and forwarded towards the Lawful Intercept service function in order to send this traffic towards the external Government Agency.
Business Drivers	Lawful Intercept can be considered as a service offered to an external Government Agency and it is associated with the need to comply with legal obligations. The business value associated with Deep Packet Inspection could be potentially increased in the future with reference to new services that need to make recognition of traffic per subscriber or per application. Examples of this type of services are quota based services that present models where customers have a periodic volume quota that once consumed, requires them to buy additional admission to the same service or to have a more basic version of the service.

	Differentiations can be made based on specific applications or particular profiles and these conditions may apply or change during some hours of the day. In case the Lawful Intercept service function is needed, the replication of the traffic at BNG level must not impact any kind of traffic and must not be perceived by the users. In addition, there is a desire to use best-of-breed service-enabling devices from separate vendors, e.g. a traditional BNG (e.g. router-based), combined with Lawful Intercept and Deep Packet Inspection service functions that can be router based or implemented over dedicated devices.
Deployment Model	**The service chain(s) may include three primary service functions on the data path:** 1. Regular BNG (L3) function 2. Lawful Intercept and Deep Packet inspection with traffic replicated at BNG level But when Lawful Intercept is not needed: 1. Regular BNG (L3) function 2. Deep Packet inspection If t Deep Packet Inspection and Lawful Intercept are not needed: 1. Regular BNG (L3) function **Geographical distribution**: it is assumed that physical systems hosting all three service functions are collocated in a single facility (e.g. in a Metropolitan Point-Of-Presence). **Administrative boundaries**: all service functions are contained within a single MBN "+" network, operated by a single service provider and administrative entity.
Actors	Multiservice Broadband Network Broadband users (residential and business customers) Broadband Network and Service Provider External Government Agency

High-level architectural context	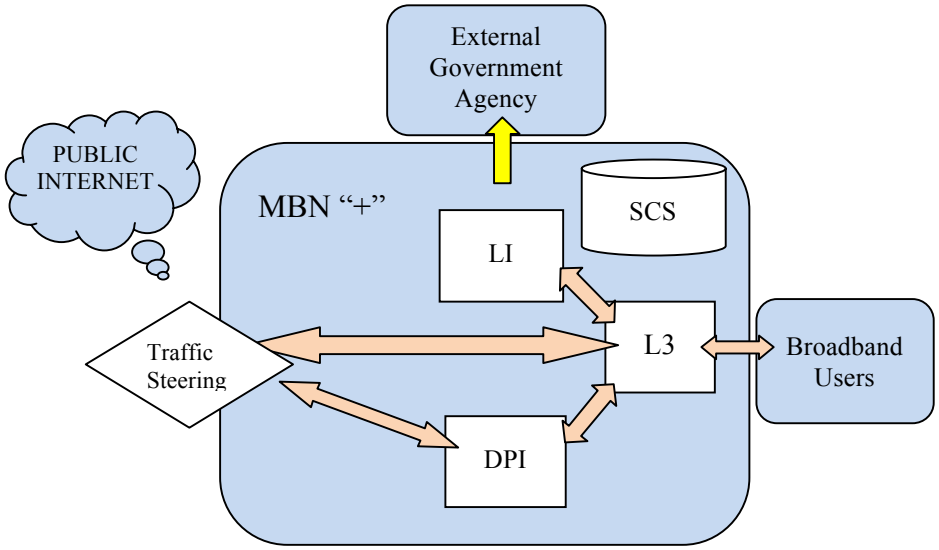 All service functions are deployed as part of an extended Multiservice Broadband Network ("MBN +"), allowing user traffic to flow to/from the public Internet, while being appropriately serviced. The BNG is configured to give L3 VPN or IP connectivity to broadband users. The chaining functions should allow the concatenation of local and remote service enabling functions. Two types of service chains are used, either BNG to DPI or BNG only, in both cases with possible traffic replication at BNG level if Lawful Intercept is needed. The traffic replication at BNG level (in case the Lawful Intercept service function is needed) can happen through different levels of granularity (e.g. all the traffic of a BNG or traffic from a whole port of the BNG or only the traffic of a subscriber), according to the appropriate local government laws. Depending on the exact terms of the service being subscribed to, the BNG may enforce various types of QoS/Policy profiles. The Lawful Intercept service function sends the traffic received by the BNG towards the external Government Agency while the DPI service function performs traffic analysis with profile-based and/or service-based differentiation. Architectural attributes (functional view): **Service chain shape**: chain of 1or 2 systems. **Performance**: no traffic impacted by replication at BNG level; Lawful Intercept must not be perceived by the users; SLAs similar to typical residential customer services or business customer services **Load balancing & Resiliency**: automated balancing between a farm of servers supporting the Deep Packet Inspection service function for a given BNG, based on simple criteria (e.g. max number of user sessions per server). **Automation & Lifecycle:** choice of exact service profile per user fully driven by an external system (e.g. AAA or SNMP), with minimum OSS integration burden and changes compared to a regular BNG-only service deployment. At the BNG

	level, traffic per user session, must be appropriately filtered and classified for steering on the proper service chain. Such filters may be dynamic and have a different lifecycle with reference to the user session (e.g. PPP, DHCP) as the need to make the traffic traverse the Deep Packet Inspection service function or to duplicate it for the Lawful Interception could be relative to a period of time shorter than the session lifecycle. **Traffic engineering**: **Symmetrical forwarding**: required for all traffic, including the one that traverses Deep Packet Inspection service function. It is important to notice that a traffic steering mechanism is necessary to make Deep Packet Inspection service function reachable for the traffic coming from the network and going towards the customers. **Other considerations (e.g. mobility, etc)**: none
Related and Derivative Use Cases	**Existing use cases**: no existing SD-326 use case appears to be directly related to this one (so far). **Use case evolution**: the use case described in this table might morph into something slightly different over time, e.g.: • The Lawful Intercept and the Deep Packet inspection functions might be located in different Points-Of-Presence • The Lawful Intercept and the Deep Packet inspection functions might be virtualized and implemented on a virtual machine
Issue(s) Spotlight	The need for symmetrical forwarding for this use case implies that the downstream traffic (from the network to the user) traverses Deep Packet Inspection service function when needed. A traffic steering mechanism, that makes the Deep Packet Inspection Service Function reachable for the traffic coming from the network and going towards the customers, is therefore necessary.
Inputs to market requirements	This use case illustrates several key needs for service chaining: - The need for symmetrical forwarding via some form of traffic steering - The first Service Function (BNG) needs to be capable of classifying traffic (including for a given user session) before steering to the appropriate service chain - The association between traffic associated with a given user session and a given service profile needs to be somehow propagated between the first Service Function (BNG) and the other Service Functions - Load balancing and steering of traffic for a given user session on a given service chain needs to be deterministic enough, so that all traffic for a given user session always traverses the same instance of Service Functions.

5. Geographically distributed Service Functions: Lawful Intercept, Deep Packet Inspection and URL Filtering

This use case is directly related to the use case "Lawful Intercept and Deep Packet Inspection" as described in 4, extending its scope to URL Filtering and geographical distribution. The related use case will be referred to as "[LI-DPI]" in the table below.

Title	Service Chaining for broadband users with geographically distributed Service Functions: Lawful Intercept, Deep Packet Inspection and URL Filtering
Service Model and Story Highlights	The Service Model as perceived by broadband users is VoIP, Video and Internet access and the connectivity is IP-based. Internet traffic may optionally be subject to URL Filtering (e.g. parental control). Inside the network, three Service Functions will be involved in addition to the traditional BNG: Lawful Intercept (LI), Deep Packet Inspection (DPI), URL Filtering. Such functions may be geographically distributed (LI and DPI at a Metropolitan POP level; URL Filtering in a remote Data Center). Lawful Intercept and Deep Packet Inspection considerations are the same as [LI-DPI]. URL Filtering is offered to customers who want to prevent their family (e.g. parental control) from accessing "inappropriate" Web sites through different profiles of filters being available, including no filtering at all. In order to improve performance while reducing cost of operation, only the traffic of users having subscribed to URL Filtering terms of service is backhauled to a remote Data Center, and passes through the URL Filtering function. The traffic of users selecting a Web filtering "null" profile and the non-HTTP traffic (notably latency-sensitive traffic like VoIP) do not pass through the URL Filtering function.
Business Drivers	Business drivers related to Lawful Intercept and Deep Packet are the same as [LI-DPI]. The business value of the URL Filtering service is to satisfy a class of users who may want to prevent their family from accessing Web sites with inappropriate content. The choice of geographical distribution is driven by a balance between data path optimization and Capex/Opex considerations. It is expected that a large percentage of traffic will be subject to Deep Packet Inspection, while a small percentage of traffic will be subject to URL Filtering. It is therefore desirable to distribute the DPI functionality at the Metropolitan level, while deploying the URL Filtering functionality in a more centralized data center. Similar to [LI-DPI], there is a desire to use best-of-breed service-enabling devices from separate vendors, e.g. a traditional BNG (e.g. router-based), combined with Lawful Intercept, Deep Packet Inspection and URL Filtering service functions that

	can be router based or implemented over dedicated devices. For the URL Filtering service function, its location in a Data Center makes it more desirable to be implemented on generic off-the-shelf servers, and to share corresponding physical resources with other applications through the use of Virtual Machines.
Deployment Model	**The service chain(s) may include four primary service functions on the data path:** 1. Regular BNG (L3) function 2. Lawful Intercept 3. URL Filtering 4. Deep Packet Inspection When Lawful Intercept is not needed, the service chain may include only three service functions: 1. Regular BNG (L3) function 2. URL Filtering 3. Deep Packet inspection For non-HTTP traffic or for service profiles without URL filtering, in case DPI is still necessary: 1. Regular BNG (L3) function 2. Deep Packet inspection If URL Filtering, Deep Packet Inspection and Lawful Intercept are all not needed: 1. Regular BNG (L3) function **Geographical distribution**: it is assumed that physical systems hosting Regular BNG (L3) function, Lawful Intercept and Deep Packet Inspection functions are collocated in a single facility, in a Metropolitan Point-Of-Presence, while the URL Filtering service function is located in a remote Data Center. **Administrative boundaries**: all service functions are contained within a single MBN "+" network that is operated by a single service provider and administrative entity. The MBN "+" network encompasses Metropolitan POPs, an IP routing backbone and Data Centers.
Actors	Multiservice Broadband Network (Metropolitan PoPs, Backbone and Data Centers) Broadband users (residential and business customers) Broadband Network and Service Provider External Government Agency

High-level architectural context

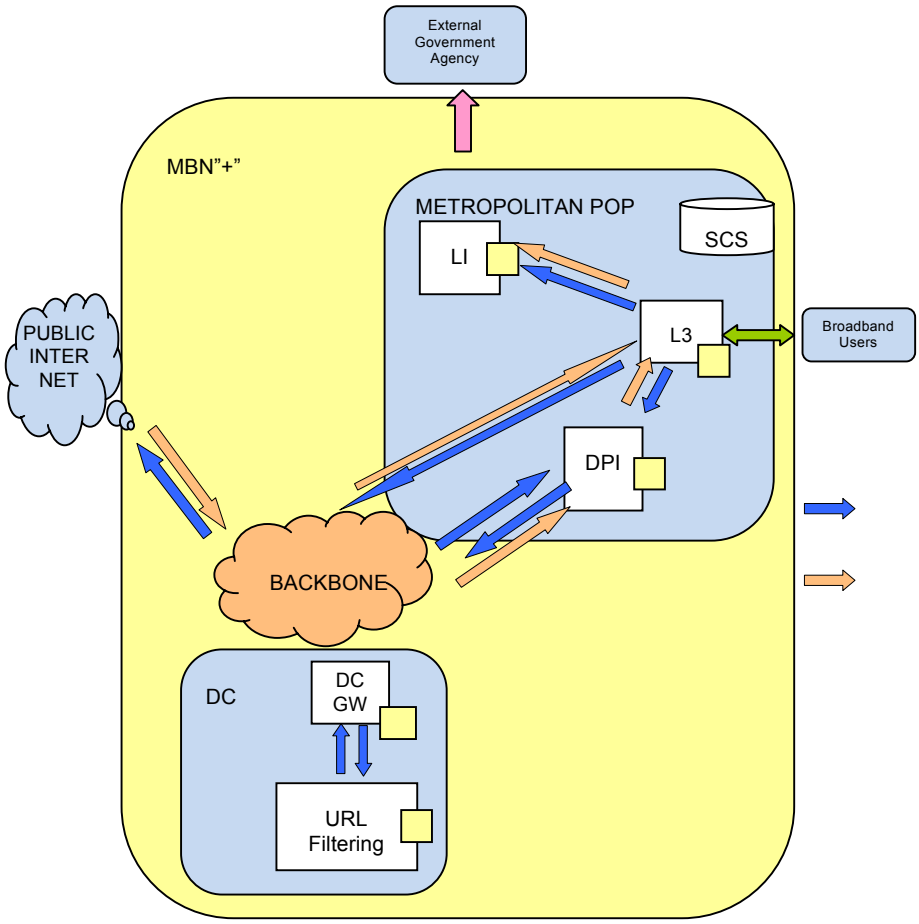

All service functions are deployed as part of an extended Multiservice Broadband Network ("MBN +"), including multiple Metropolitan PoPs and a remote Data Center interconnected by an IP backbone, allowing user traffic to flow to/from the public Internet, while being appropriately serviced.

The BNG is configured to provide L3 VPN or Internet connectivity to broadband users. The chaining functions must allow the concatenation of local (e.g. Metro-POP) and remote (e.g. Data Center) service functions.

Four main types of service chains are used: BNG to URL Filtering and DPI, BNG to URL Filtering, BNG to DPI or BNG only, in all cases with possible traffic replication at BNG level if Lawful Intercept is needed.

The following table describes the different possible chains for the upstream and the downstream packets.

DPI	**LI**	**URL FILTERING**	
NO	NO	NO	**UP**: 1-Public Internet **DW**: Public Internet-1

NO	YES	NO	**UP**: 1-Public Internet + 1-3 **DW**: Public Internet-1+1-3
NO	NO	YES	**UP**: 1-4-5-4-Public Internet **DW**: Public Internet-1
NO	YES	YES	**UP**: 1-4-5-4-Public Internet + 1-3 **DW**: Public Internet-1+ 1-3
YES	NO	NO	**UP**: 1-2-Public Internet **DW**: Public Internet-2-1
YES	YES	NO	**UP**: 1-2- Public Internet + 1-3 **DW**: Public Internet-2-1+ 1-3
YES	NO	YES	**UP**: 1-4-5-4-2-Public Internet **DW**: Public Internet-2-1
YES	YES	YES	**UP**: 1-4-5-4-2-Public Internet + 1-3 **DW**: Public Internet-2-1 + 1-3

Lawful Interception works as described in [LI-DPI], with selective traffic replication at BNG level, and the Lawful Intercept service function sending traffic received by the BNG towards the external Government Agency. Deep Packet Inspection works as described in [LI-DPI], subject to the considerations described below.

If the URL Filtering function is needed, traffic is steered to the DC by some means (e.g. through the use of tunnels - multiple possible architectural approaches could be envisioned). A DC Gateway is located at the entry point of the Data Center. This DC Gateway receives through the backbone user traffic from the BNG and forwards it towards the URL filtering service function. The URL Filtering function is typically implemented as a farm of servers/VMs, hence requiring a form of load balancing; it filters traffic according to URLs and service profiles, then forwards it towards the Backbone to make it reach the DPI service function (if needed), or directly the Public Internet. Downstream traffic from the Public Internet does not traverse the URL Filtering service function, and should entirely bypass Data Centers. In case URL Filtering service function and DPI service function are both needed, the URL Filtering service function must be traversed first in order to prevent non-relevant traffic to be accounted for (e.g. if the URL requested is not allowed).

Architectural attributes (functional view):

Service chain shape: chain of 1, 2 or 3 service functions. The chain may span several physical locations. It may be ‘closed’ (a loop back to the BNG) or open (two entry points, upstream and downstream), or even geographically open (entry points and end points in distinct physical locations).

Performance: no traffic impacted by replication at BNG level; low-latency traffic (e.g. VoIP) not impacted by URL Filtering; Lawful Intercept must not be

	perceived by the users; SLAs similar to typical residential customer services or business customer services. **Load balancing & Resiliency**: load-balancing needs to be deterministic, keeping and enforcing affinity between flows associated with a given user session for a given service chain (e.g. traverse the same set of Service Functions). In addition, the typical use of a farm of servers to support the URL Filtering service function in a Data Center will need to be accounted for, with an appropriate form of load-balancing when related traffic enters the Data Center. **Automation & Lifecycle:** same as [LI-DPI]. **Traffic engineering**: some level of Traffic Engineering may be required to provide minimal resource guarantees on the (tunneled) data path(s) supporting Metro-POP to remote-DC connectivity. **Symmetrical forwarding**: required for traffic that traverses Deep Packet Inspection service function, as described in [LI-DPI]. Symmetrical forwarding is NOT required (and actually undesirable) for traffic that traverses the URL Filtering function, leading to highly asymmetrical data paths through the MBN network and the backbone. **Other considerations (e.g. mobility, etc)**: none.
Related and Derivative Use Cases	**Related use cases**: this use case is directly related to the use case "Lawful Intercept and Deep Packet Inspection" as described in [LI-DPI], extending its scope to URL Filtering and geographical distribution.
Issue(s) Spotlight	Symmetrical forwarding is necessary for the traffic that traverses the Deep Packet Inspection service function. Asymmetrical forwarding is highly desired for traffic that traverses URL Filtering service function, which only applies to (upstream) traffic originated by the users. The service chains may span multiple physical locations, which are geographically distributed. This implies relatively complex service chains, as exemplified by the following picture. This shows the case where URL filtering and Deep Packet Inspection service functions are both required. For the upstream (from the end user) direction, the service chain goes from one location (Metro-POP) to another (remote Data Center), and comes back to the first location (Metro-POP). For the downstream direction (towards the user), the service chain is much simpler and is contained in the Metro-POP. (In the picture the blue line refers to the upstream direction and the orange to the downstream one).

	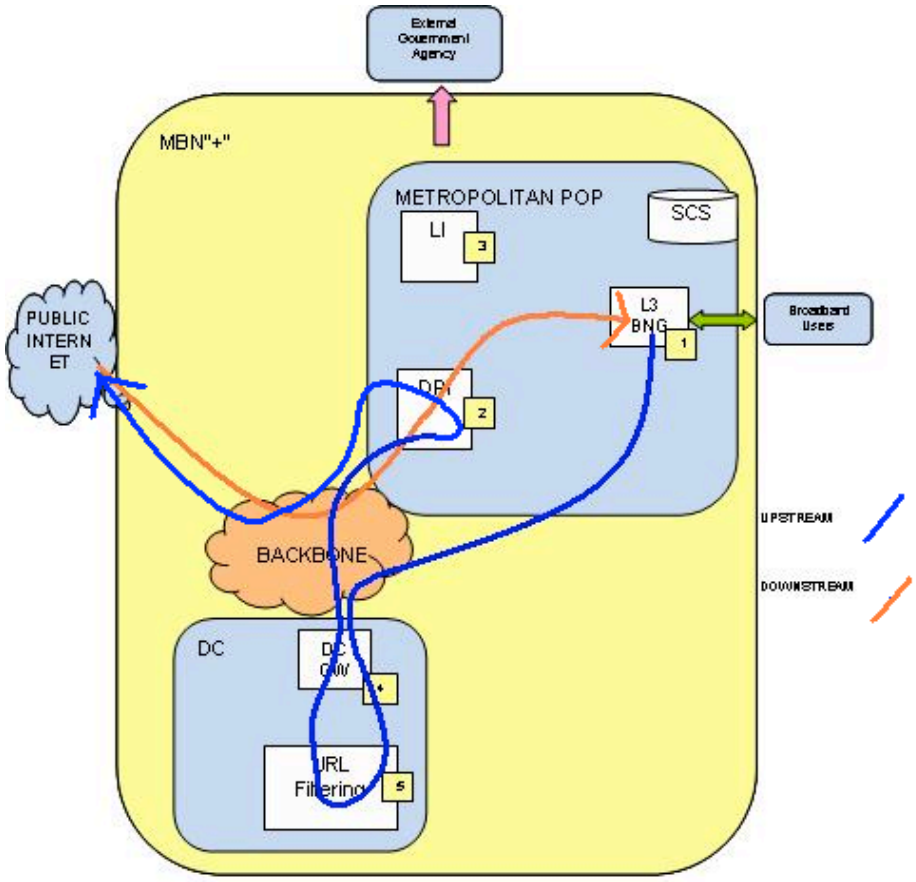 Some advanced form of tunneling and traffic steering mechanisms is therefore necessary, spanning routing-centric network infrastructure (Metro-POP and backbone) as well as switching-centric network infrastructure (Data Center), and interconnecting tunneling 'overlay' technologies as appropriate through a Data Center Gateway of sorts, while enforcing proper sequencing of Service Functions.
Inputs to market requirements	In addition to the key needs expressed in [LI-DPI] (e.g. traffic classification, traffic steering and symmetrical forwarding), this use case illustrates several additional key needs for service chaining: - The association between traffic with a given user session and a given service profile needs to be somehow propagated between the first Service Function (BNG) and other Service Functions that can either be collocated with the BNG or be hosted by a remote site (e.g. Data Center). - Load balancing and steering of traffic for a given user session on a given service chain needs to be deterministic (as described in [LI-DPI]) and to support a stage of load-balancing at the entry point of the Data Center, to accommodate a farm of (virtual) servers supporting a Service Function (e.g. URL Filtering). - Flexible and constrained tunneling schemes are needed to support service chains that may span distributed locations, including a possible 'U-turn'

	(e.g. Metro-POP to DC to Metro-POP). - Flexible tunnel interconnections are needed to adequately stitch together forms of tunnels appropriate for routing-centric network infrastructure (e.g. Metro-POP and backbone) and switching-centric network infrastructure (e.g. Data Center). - Asymmetrical forwarding needs to be supported for traffic that traverses Service Functions hosted by a remote Data Center (e.g. URL Filtering) via some form of traffic steering and tunneling mechanism, while the return path should entirely bypass such Data Center facility.