Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 23, 2015

P. Quinn, Ed.
Cisco Systems, Inc.
T. Nadeau, Ed.
Brocade
February 19, 2015

Service Function Chaining Problem Statement
draft-ietf-sfc-problem-statement-13.txt

Abstract

This document provides an overview of the issues associated with the deployment of service functions (such as firewalls, load balancers, etc.) in large-scale environments. The term service function chaining is used to describe the definition and instantiation of an ordered list of instances of such service functions, and the subsequent "steering" of traffic flows through those service functions.

The set of enabled service function chains reflect operator service offerings and is designed in conjunction with application delivery and service and network policy.

This document also identifies several key areas that the SFC working group will investigate to guide its architectural and protocol work and associated drafts.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 23, 2015.

Table of Contents

1. Introduction

The delivery of end-to-end services often require various service functions including traditional network service functions (for example firewalls and server load balancers), as well as application-specific features such as http header manipulation. Service functions may be delivered within the context of an isolated user (e.g. a tenant), or shared amongst many users/user groups.

Current service function deployment models are often tightly coupled to network topology and physical resources resulting in relatively rigid and static deployments. The static nature of such deployments greatly reduces, and in many cases, limits the ability of an operator to introduce new or modify existing services and/or service functions. Furthermore there is a cascading effect: changing one (or more) elements of a service function chain often affects other elements in the chain and/or the network elements used to construct the chain.

This issue is particular acute in elastic service environments that require relatively rapid creation, destruction or movement of physical or virtual service functions or network elements. Additionally, the transition to virtual platforms requires an agile service insertion model that supports elastic and very granular service delivery, post-facto modification and the movement of service functions and application workloads in the existing network. The service insertion model must also retain the network and service policies and the ability to easily bind service policy to granular information such as per-subscriber state.

This document outlines the problems encountered with existing service deployment models for Service Function Chaining (SFC) (often referred to simply as service chaining (in this document the terms will be used interchangeably), as well as the problems of service chain creation, deletion, modification/update, policy integration with service chains, and policy enforcement within the network infrastructure. The document highlights three key areas of WG focus for addressing the issues highlighted in this draft that will form the basis for the possible WG solutions that address the current problems.

1.1. Definition of Terms

Classification: Locally instantiated matching of traffic flows against policy for subsequent application of the required set of network service functions. The policy may be customer/network/service specific.

Network Overlay: A logical network built, via virtual links or packet encapsulation, over an existing network (the underlay).

Network Service: An offering provided by an operator that is delivered using one or more service functions. This may also be referred to as a composite service. The term "service" is used to denote a "network service" in the context of this document.

Note: Beyond this document, the term "service" is overloaded with varying definitions. For example, to some a service is an offering composed of several elements within the operator's network, whereas for others a service, or more specifically a network service, is a discrete element such as a "firewall". Traditionally, such services (in the latter sense) host a set of service functions and have a network locator where the service is hosted.

Service Function: A function that is responsible for specific treatment of received packets. A Service Function can act at various layers of a protocol stack (e.g., at the network layer or other OSI layers). As a logical component, a Service Function can be realized as a virtual element or be embedded in a physical network element. One or more Service Functions can be embedded in the same network element. Multiple occurrences of the Service Function can exist in the same administrative domain.

A non-exhaustive list of service functions includes: firewalls, WAN and application acceleration, Deep Packet Inspection (DPI), server load balancers, NAT44 [RFC3022], NAT64 [RFC6146], HTTP Header Enrichment functions, TCP optimizer.

The generic term "L4-L7 services" is often used to describe many service functions.

Service Function Chain (SFC): A service function chain defines an ordered or partially ordered set of abstract service functions (SFs) and ordering constraints that must be applied to packets and/or frames and/or flows selected as a result of classification. An example of an abstract service function is "a firewall". The implied order may not be a linear progression as the architecture allows for SFCs that copy to more than one branch, and also allows for cases where there is flexibility in the order in which service functions need to be applied. The term service chain is often used as shorthand for service function chain.

Service Overlay: An overlay network created for the purpose of forwarding data to required service functions.

Service Topology: The service overlay connectivity forms a service topology.

2. Problem Space

The following points describe aspects of existing service deployments that are problematic, and that the Service Function Chaining (SFC) working group aims to address.

2.1. Topological Dependencies

Network service deployments are often coupled to network topology, whether it be physical or virtualized, or a hybrid of the two. For example, use of a firewall requires that traffic flow through the firewall, which require means placing the firewall on the network path (often via creation of VLANs), or architecting the network topology to steer traffic through the firewall. Such dependency imposes constraints on service delivery, potentially inhibiting the network operator from optimally utilizing service resources, and reduces flexibility. This limits scale, capacity, and redundancy across network resources.

These topologies serve only to "insert" the service function (i.e., ensure that traffic traverses a service function); they are not required from a native packet delivery perspective. For example, firewalls often require an "in" and "out" layer-2 segment and adding a new firewall requires changing the topology (i.e., adding new layer-2 segments and/or IP subnets).

As more service functions are required - often with strict ordering - topology changes are needed in "front" and "behind" each service function resulting in complex network changes and device configuration. In such topologies, all traffic, whether a service function needs to be applied or not, often passes through the same strict order.

The topological coupling limits placement and selection of service functions: service functions are "fixed" in place by topology and therefore placement and service function selection taking into account network topology information such as load, new links, or traffic engineering is often not possible.

A common example is web servers using a server load balancer as the default gateway. When the web service responds to non-load balanced traffic (e.g., administrative or backup operations) all traffic from the server must traverse the load balancer forcing network administrators to create complex routing schemes or create additional interfaces to provide an alternate topology.

2.2. Configuration complexity

A direct consequence of topological dependencies is the complexity of the entire configuration, specifically in deploying service function chains. Simple actions such as changing the order of the service functions in a service function chain require changes to the logical and/or physical topology. However, network operators are hesitant to make changes to the network once services are installed, configured and deployed in production environments for fear of misconfiguration and consequent downtime. All of this leads to very static service delivery deployments. Furthermore, the speed at which these topological changes can be made is not rapid or dynamic enough as it often requires manual intervention, or use of slow provisioning systems.

2.3. Constrained High Availability

Since traffic reaches many service functions based on network topology, alternate, or redundant service functions must be placed in the same topology as the primary service.

An effect of topological dependency is constrained service function high availability. Worse, when modified, inadvertent non-high availability or downtime can result.

2.4. Consistent Ordering of Service Functions

Service functions are typically independent; service function_1 (SF1)...service function_n (SFn) are unrelated and there is no notion at the service layer that SF1 occurs before SF2. However, to an administrator many service functions have a strict ordering that must be in place, yet the administrator has no consistent way to impose and verify the ordering of the service functions that are used to deliver a given service. Furthermore, altering the order of a deployed chain is complex and cumbersome.

2.5. Application of Service Policy

Service functions rely on topology information such as VLANs or packet (re)classification to determine service policy selection, i.e., the service function specific action taken. Topology information is increasingly less viable due to scaling, tenancy and complexity reasons. Topology-centric information often does not convey adequate information to the service functions, forcing functions to individually perform more granular classification. In other words, the topology information is not granular enough, and its semantics often overloaded.

2.6. Transport Dependence

Service functions can and will be deployed in networks with a range of network transports, including network under and overlays, such as Ethernet, GRE, VXLAN, MPLS, etc. The coupling of service functions to topology may require service functions to support many transport encapsulations or for a transport gateway function to be present.

2.7. Elastic Service Delivery

Given that the current state of the art for adding/removing service functions largely centers around VLANs and routing changes, rapid changes to the deployed service capacity (increasing or decreasing) can be hard to realize due to the risk and complexity of VLANs and/or routing modifications.

2.8. Traffic Selection Criteria

Traffic selection is coarse, that is, all traffic on a particular segment traverses all service functions whether the traffic requires service enforcement or not. This lack of traffic selection is largely due to the topological nature of service deployment since the forwarding topology dictates how (and what) data traverses which service function(s). In some deployments, more granular traffic selection is achieved using policy routing or access control filtering. This results in operationally complex configurations and is still relatively coarse and inflexible.

2.9. Limited End-to-End Service Visibility

Troubleshooting service related issues is a complex process that involve both network-specific and service-specific expertise. This is especially the case when service function chains span multiple DCs, or across administrative boundaries. Furthermore, the physical and virtual environments (network and service), can be highly divergent in terms of topology and that topological variance adds to these challenges.

2.10. Per-Service Function (re)Classification

Classification occurs at each service function independent from previously applied service functions since there are limited mechanisms to share the detailed classification information between services. The classification functionality often differs between service functions, and service functions may not leverage the classification results from other service functions.

2.11. Symmetric Traffic Flows

Service function chains may be unidirectional or bidirectional depending on the state requirements of the service functions. In a unidirectional chain traffic is passed through a set of service functions in one forwarding direction only. Bidirectional chains require traffic to be passed through a set of service functions in both forwarding directions. Many common service functions such as DPI and firewall often require bidirectional chaining in order to ensure flow state is consistent.

Existing service deployment models provide a static approach to realizing forward and reverse service function chain association most often requiring complex configuration of each network device throughout the SFC. In other words, the same complex network configuration must be in place for both "directions" of the traffic, effectively doubling the configuration and associated testing. Further, if partial symmetry is required (i.e. only some of the services in the chain required symmetry), the network configuration complexity increases since the operator must ensure that the exceptions -- the services that do not need the symmetry flow -- are handled correctly via unique configuration to account for their requirements.

2.12. Multi-vendor Service Functions

Deploying service functions from multiple vendors often require per-vendor expertise: insertion models differ, there are limited common attributes and inter-vendor service functions do not share information, hence the need for standards to ensure interoperability.

3. Service Function Chaining

Service Function Chaining aims to address the aforementioned problems associated with service deployment. Concretely, the SFC working group will investigate solutions that address the following elements:

3.1. Service Overlay

Service function chaining utilizes a service specific overlay that creates the service topology. The service overlay provides service function connectivity, built "on top" of the existing network topology and allows operators to use whatever overlay or underlay they prefer to create a path between service functions, and to locate service functions in the network as needed.

Within the service topology, service functions can be viewed as resources for consumption and an arbitrary topology constructed to connect those resources in a required order. Adding new service functions to the topology is easily accomplished, and no underlying network changes are required.

Lastly, the service overlay can provide service specific information needed for troubleshooting service related issues.

3.2. Service Classification

Classification is used to select which traffic enters a service overlay. The granularity of the classification varies based on device capabilities, customer requirements, and services offered. Initial classification determines the service function chain required to process the traffic. Subsequent classification can be used within a given service function chain to alter the sequence of service functions applied. Symmetric classification ensures that forward and reverse chains are in place. Similarly, asymmetric -- relative to required service function -- chains can be achieved via service classification.

3.3. SFC Encapsulation

The SFC encapsulation enables the creation of a service chain in the data plane and can convey information about the chain such as chain identification and OAM status.

The SFC encapsulation also carries data plane metadata which provides the ability to exchange information between logical classification points and service functions (and vice versa) and between service functions. Metadata is not used as forwarding information to deliver packets along the service overlay.

Metadata can include the result of antecedent classification and/or information from external sources. Service functions utilize metadata, as required, for localized policy decisions.

In addition to sharing of information, the use of metadata addresses several of the issues raised in section 2, most notably by decoupling policy from the network topology, and by removing the need for per-service function classification (and re-classification) described in section 2.10.

A common approach to service metadata creates a common foundation for interoperability between service functions, regardless of vendor.

4. IANA Considerations

This document makes no request to IANA.

5. Security Considerations

Although this problem statement does not introduce any protocols, when considering service function chaining, the three main areas begin investigated (see section 3) by the WG have security aspects that warrant consideration.

Service Overlay: The service overlay will be constructed using existing transport protocols (e.g. MPLS, VXLAN) and as such is subject to the security specifics of the transport selected. If an operator requires authenticity and/or confidentiality in the service overlay, a transport (e.g. IPSec) that provides such functionally can be used.

Classification: Since classification is used to select the appropriate service overlay, and required service encapsulation details, classification policy must be both accurate and trusted. Conveying the policy to a SFC-edge device node may be done via a multitude of methods depending on an operator's existing provisioning practices and security posture.

Additionally, traffic entering the SFC domain and being classified may be encrypted thus limiting the granularity of classification. The use of pervasive encryption varies based on type of traffic, environment and level of operator control. For instance a large enterprise can mandate how encryption is used by its users, whereas a broadband provider likely does not have the ability to do so.

The use of encrypted traffic however does not obviate the need for SFC (nor the problems associated with current deployment models described herein), rather when encrypted traffic must be classified, the granularity of such classification must adapt. In such cases, service overlay selection might occur, for example, using outer (i.e. unencrypted) header information, on the presence of encryption, or via external information about the packets.

SFC Encapsulation: As described in section 3, the SFC encapsulation carries information about the SFC, and data plane metadata. Depending on environment and security posture, the SFC encapsulation might need to be authenticated and/or encrypted. The use of an appropriate overlay transport as described above can provide data plane confidentially and authenticity.

The exchange of SFC encapsulation data such as metadata must originate from trusted source(s) and, if needed, be subject to authenticity and confidentiality during the exchange to the various SFC nodes.

SFC and Multi-tenancy: If tenant isolation is required in an SFC deployment, an appropriate network transport overlay that provides adequate isolation and identification can be used. Additionally, tenancy might be used in the selection of the appropriate service chain, however, as stated, the network overlay is still required to provide transport isolation. SF deployment and how specific SFs might or might not be allocated per tenant is outside the scope of this document.

The SFC Architecture draft present a more complete review of the security implications of a complete SFC architecture.

6. Contributors

The following people are active contributors to this document and have provided review, content and concepts (listed alphabetically by surname):

Puneet Agarwal
Broadcom
Email: pagarwal@broadcom.com

Mohamed Boucadair
France Telecom
Email: mohamed.boucadair@orange.com

Abhishek Chauhan
Citrix
Email: Abhishek.Chauhan@citrix.com

Uri Elzur
Intel
Email: uri.elzur@intel.com

Kevin Glavin
Riverbed
Email: Kevin.Glavin@riverbed.com

Ken Gray
Cisco Systems
Email: kegray@cisco.com

Jim Guichard
Cisco Systems
Email:jguichar@cisco.com

Christian Jacquenet
France Telecom
Email: christian.jacquenet@orange.com

Surendra Kumar
Cisco Systems
Email: smkumar@cisco.com

Nic Leymann
Deutsche Telekom
Email: n.leymann@telekom.de

Darrel Lewis
Cisco Systems

Email: darlewis@cisco.com

Rajeev Manur
Broadcom
Email:rmanur@broadcom.com

Brad McConnell
Rackspace
Email: bmcconne@rackspace.com

Carlos Pignataro
Cisco Systems
Email: cpignata@cisco.com

Michael Smith
Cisco Systems
Email: michsmit@cisco.com

Navindra Yadav
Cisco Systems
Email: nyadav@cisco.com

7. Acknowledgments

The authors would like to thank David Ward, Rex Fernando, David McDysan, Jamal Hadi Salim, Charles Perkins, Andre Beliveau, Joel Halpern and Jim French for their reviews and comments.

Additionally, the authors would like to thank the IESG and Benjamin Kaduk for their detailed reviews and suggestions.

8. Informative References

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, January 2001.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

Authors' Addresses

Paul Quinn (editor)
Cisco Systems, Inc.

Email: paulq@cisco.com

Thomas Nadeau (editor)
Brocade

Email: tnadeau@lucidvision.com