Nework working group
Internet Draft
Intended status: Informational
Expires: January 2014

L. Dunbar
D. Eastlake
Huawei

July 11, 2013

Layer 4-7 Service Chain problem statement
draft-dunbar-l4-l7-sc-problem-statement-00.txt

Status of this Memo

Copyright Notice

Abstract

This draft analyzes the taxonomy of Layer 4-7 Services and gives two examples of Layer 4-7 service chain, one from a traffic steering perspective and another one from a Layer 7 perspective. The intent is to emphasize their unique issues and challenges.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC-2119 0.

The term "traffic steering" and "traffic forwarding" are used interchangeably in this draft.

Table of Contents

1. Introduction

This draft analyzes the taxonomy of Layer 4-7 Services and gives two examples of Layer 4-7 service chain, one from a traffic steering perspective and another one from a Layer 7 perspective. The intent is to emphasize their unique issues and challenges.

Layer 4-7 services and service chains have been discussed in many forums, such as ETSI NFV, ONF, and IETF I2RS Interim meetings. However, people from different background frequently have different interpretations of Layer 4-7 services and service chains. For example, network vendors tend to view "Layer 4-7 Service Chains" as forwarding (or steering) traffic to a sequence of service modules based on Layer 4-7 fields, whereas Layer 4-7 vendors may view "service chains" as reassembling whole HTTP messages (which could be in multiple data frames) and applying the needed functions (e.g. Content Optimization or App Security) based on some logics formulated from the message content. This draft starts with analyzing the taxonomy Layer 4-7 services and service chains.

2. Terminology

DPI	Deep Packet Inspection
FW	Firewall

3. Taxonomy of Layer 4-7 Services

Layer 4-7 Services can be broadly broken into two categories:

1) Layer 4-7 Traffic Steering: a functional module in a device that forwards data packets received from one port to another port based on Layer 4 to Layer 7 fields in the data packets.

2) Layer 4-7 Service Function: a functional module that performs Layer 4 to 7 functions, such as Firewall, DPI, TCP accelerator, NAT, etc. When Layer 4-7 service function is instantiated on a standalone physical or virtual device, it is called Layer 4-7 Service module throughout this draft. Layer 4-7 functions can

also be embedded in another device, such as router/switch or other devices.

3.1. Layer 4-7 Traffic Steering (or Forwarding)

Layer 4-7 Traffic Steering (or forwarding) basically forwards data packets received from one port to another port based on some higher layer fields in the data packets.

There are multiple types of traffic steering:

- Fixed header based forwarding: traffic steering based on header fields that have fixed position in the data packets:
  - Forwarding based on Layer 2-3 header fields, such as MAC or IP Destination Address, MPLS label, or VLAN ID.
  - Forwarding based on Layer 4 header (TCP or UDP).
  - QoS header based forwarding.
- Layer 7 based forwarding: traffic steering (or forwarding) based on the payload (L7) of data packets.

  Multiple data packets may carry some meaningful data, like one HTTP message. Under this scenario, multiple data packets have to be examined before meaningful data can be extracted for making Layer 7 based forwarding decision.

Since routers/switches all forward data packets based Layer 2 or 3 header, for ease of description "Service Chain Steering Point (or Node)" is used throughout this draft to refer to the entities that steer traffic to a sequence of service modules.

Note: the Layer 4-7 traffic steering could also steer packets to a service module that applies non-Layer4-7 functions.

3.2. Layer 4-7 Service Function

A Layer 4-7 Service Function, or service module if it is in a standalone device or virtual device, performs a Layer 4 to 7 function based on packets received. One service module can contain multiple service functions. Examples are: Firewall, DPI,

TCP accelerator, NAT, etc. Service Module could be Proxy based or Packet Based. Note the criteria to apply Layer 4-7 functions can be based on Layer 2 or 3 fields of the data packets received. On traditional routers/switches, there are Layer 2 or 3 service functions, such as frame fragmentation and reassembly. Layer 2 or 3 service functions are out of the scope of this draft.

A Layer 7 service function can be very different from a Layer 4 service function. It is necessary to differentiate them. To be specific, there are

- Layer 4 service function
- Layer 7 service function
- Layer 4 service function with some Layer 7 intelligence.

The service modules can be further distinguished by

- Proxy based service functions: these service functions terminate original packets, may reassemble multiple packets, reopen a new connection, or formulate new packets based on the received packets.

- Packet based service functions: these service modules maintain original packets, i.e. they don't make changes to packets traversed through except possibly to metadata such as VLAN tags.

An entity (physical or virtual device) that can forward packets after one service module to another service module is considered as having two functions: a Service Function integrated with a Traffic Steering function.

3.3. Service Module connection to Service Chain Steering Points

Service modules can be connected to Service Chain steering points (such as routers/switches) in various ways:

- A service module can be embedded in a traffic steering node (i.e. embedded in a router or a switch).

In this case, the service module doesn't need an address to receive data packets. The forwarding entity can send packets that meet the steering criteria directly to the service module regardless of the destination addresses in the packets. The Service module always sends the processed packets back to the forwarding entity regardless of the destination addresses in the packets.

- A service module can be one hop away from a traffic steering node

  The one hop between the Service Chain Steering node and the service module can be a physical link (e.g. Ethernet link) or one virtual tunnel (e.g. VxLAN).

  If the one hop is a physical Ethernet link, there would be a Link Header, i.e. an outer MAC header, added to the data packets that meet the steering criteria, with MAC Source Address being the Service Chain Steering Node and MAC Destination Address being the Service module for packets from the Service Chain Steering node to the service module.

  For the reverse direction over this link, i.e. after the service module process the packets, the MAC Source Address is the Service Module and the MAC Destination Address is the Service Chain Steering node.

  The one hop link can be a transparent link, i.e. no link address is added to the data packets on the link between the Service Chain Steering node and Service Module. This scenario is considered the same as a service module being embedded in the Service Chain Steering node.

  The one hop between the Service Steering node and a service module can also be a tunnel, like a VxLAN tunnel. Under this case, the tunnel header has to be added to the data packets that meet the steering criteria for those packets to be sent to the service modules. After the service module processes the data packets, the Tunnel header has to be added to the packets for them to be sent back to the Service Chain Steering node.

- A service module can be multiple hops away from a Service Chain Steering node

4. Challenges of Service Chain from Traffic Steering Perspective

From user's perspective, the service chain is a sequence of service functions, such as Chain#1 {s1, s4, s6}, Chain#2{s4, s7} applied to a flow. A flow is loosely used in this document to refer to a selective of packets that meet certain criteria. Some users might not care at which points in the network the selected flow is steered to those service modules as long as the sequence of the service modules is correct.

From the traffic steering perspective, a Service Chain guarantees that specific data flows go through a specific sequence of service modules at designated points along the flow paths in the network, as shown in the figure below. The service modules perform some functions on the data packets in the flows, such as Firewall, NAT, QoS insertion, etc.

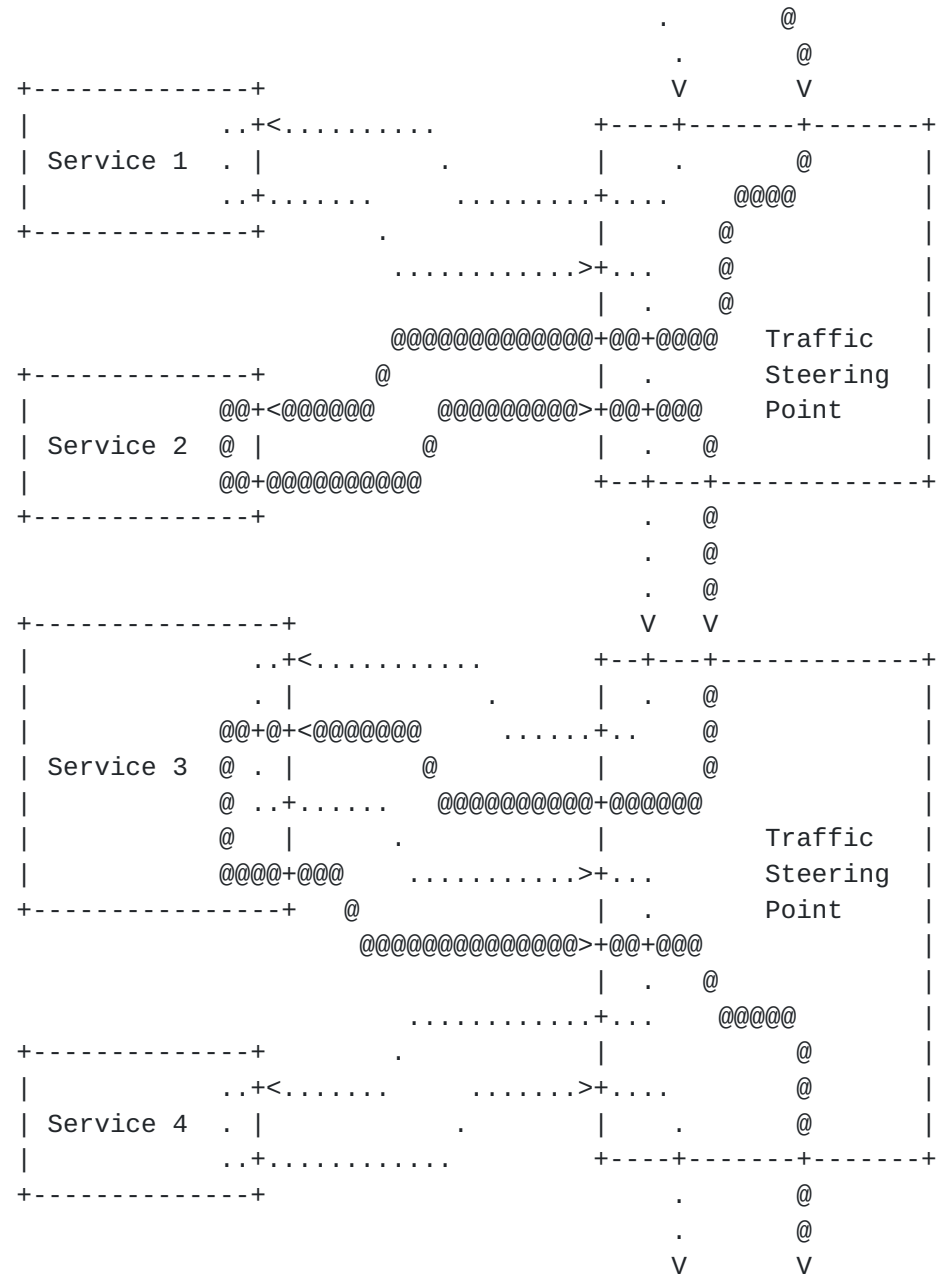

Figure 1: Simple Service Chain from Traffic Steering Point of View

4.1. Challenges of Layer 4-7 traffic Steering

Very often the criteria for steering flows to service modules are based on higher layer headers, such as TCP header, HTTP header, etc.

Most of deployed switches/routers are very efficient in forwarding packets based on Layer 2 or Layer 3 headers, such as MAC/IP destination addresses, or VLAN/MPLS labels but have limited capacity for forwarding data packets based on higher layer header. As of today, differentiating data packets based on higher layer headers depends on ACLs (Access Control List field matching) or DPI, both of which are relatively expensive and extensive use of such facilities may limit the bandwidth of switches/routers.

4.2. Challenge of traffic steering along service chain

From traffic steering point of view, one service chain consists of:

- Identifier
- {Steering point List}
- Steering Point #1, {list of Service Modules}
- Steering Point #2, {list of Service Modules}
- ?

Two service chains with the same sequence of service modules but different steering points should be considered as two different service chains from traffic steering point of view.

Some service modules change values in data packets, such as NAT changing the address fields. If any of those fields are used in traffic steering along the service chain, the criteria can be different before and after those the service modules.

Even though it is out of the scope of this draft, it is assumed that the Service Chain Orchestration System can create service chains in a way that allows each service chain to be shared by many flows while maintaining optimized utilization of network resources.

4.3. Challenges of Flow Marking for Service Chain

The policy for associating flows with their service chains can be complicated and could be dynamic. Sometimes it might not be possible to predict what traffic is traversed through and which paths traversed by.

The entity that is responsible for associating flows with their specific service Chains is called Service Chain Marking Functional Module in this document. The Service Chain Marking Functional Module can encounter flows that don't match with any policies. External entity (or controller) might be needed for a Service Chain Marking Functional Module to make appropriate decision.

Multiple flows can share one service chain. The criteria to select flows to be associated with their service chain could be different. For example, for one service chain "A" shared by Flow X, Y, Z:

- Criteria for Flow X to the Service Chain "A" are TCP port
- Criteria for Flow Y to the Service Chain "A" are Destination Address
- Criteria for Flow Z to the Service Chain "A" are MPLS label.

4.4. Ways to Minimize Impact to Existing Network

To minimize impact to deployed network elements (switches/routers), traffic flows can be classified or marked based on service chain requirement at network ingress edges, as shown in the diagram below.

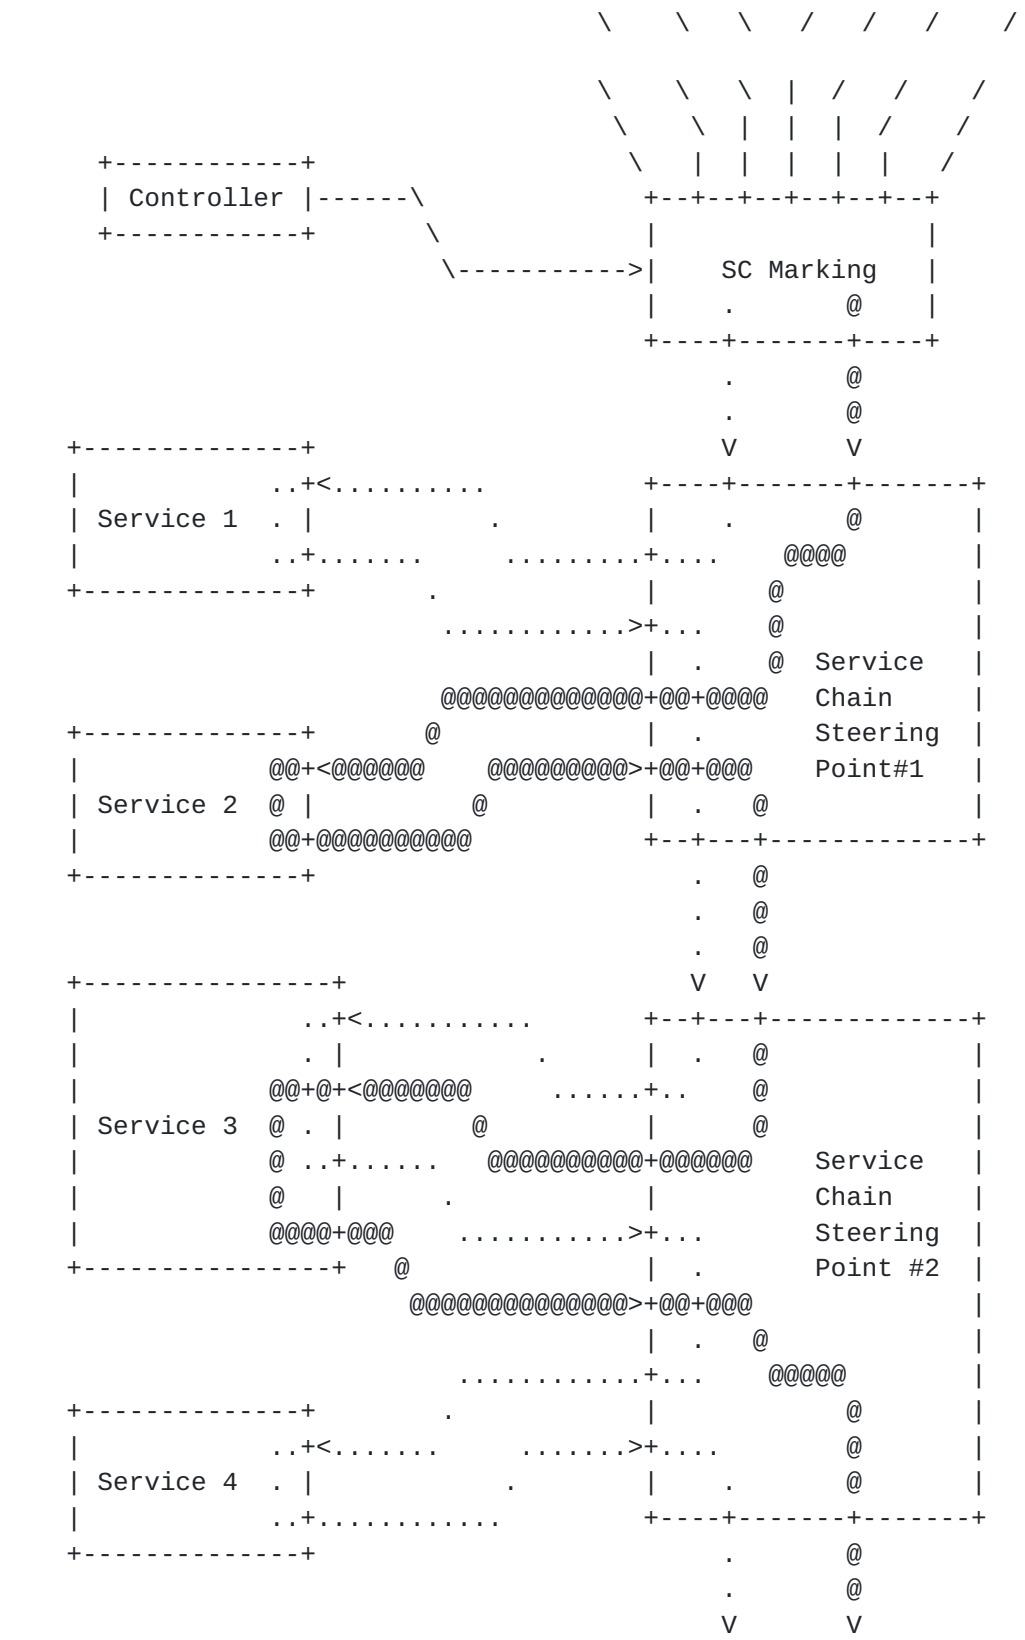

Figure 2: Service Chain Marking At Ingress

The purpose of a Service Chain Marking Functional Module is to add a unique Service Chain Label (e.g. Layer 2 or 3 Label) to the packets in the flow. Such a Layer 2 or 3 Label makes it easier for subsequent nodes along the flow path to steer the flow to the service modules specified by the flow's service chain. The network elements that have the Service Chain Marking Function are most likely network ingress edge nodes, such as Broadband Network Gateways, Cell Site Gateways, etc.

For example, the Service Chain Marking Functional Module can mark packets in a flow with a VLAN or MPLS label, based on the flow's service chain requirement.

In some situations, like service chain for wireless subscribers, many flows (i.e. subscribers) have common service chain requirements. Under those situations, the Service Chain Marking Functional Module can mark multiple flows with the same service chain requirement using the same Layer 2 or 3 Label, which effectively aggregates those flows into one service chain.

To minimize changes to deployed network elements, a small number of nodes in network can be designated to have the responsibility of steering traffic to the designated service modules. For ease of description, those nodes are called Service Chain Steering Points in this draft.

Overlay tunnels, such as VxLAN, can be used to force flows to traverse their designated Service Chain Steering Points. By using overlay tunnels, the existing network elements don't need to change any forwarding behavior.

For service chains that are shared by a great number of flows, they can be pre-provisioned. For example, if VLAN ID=10 is the service chain that need to traverse "Service-1" at Steering Point #1 and "Service-3" at Steering Point #2, the forwarding rule for VLAN ID=10 can be pre-configured at Steering Point #1 and Steering Point #2.

5. Challenge of Service Chain from the Layer 7 Perspective

From the Layer 7 perspective, the service chain can be much more complex. As shown in the figure below, the service modules to be chained can depend on the HTTP message request and reply. The service chain steering point may have to examine the whole HTTP message to determine the specific sequence of service modules for packets to traverse through. The HTTP message might have to be extracted from multiple data packets. Sometimes, the logic to steer traffic to chain of service modules might depend on the data retrieved from a database based on messages constructed from packets.

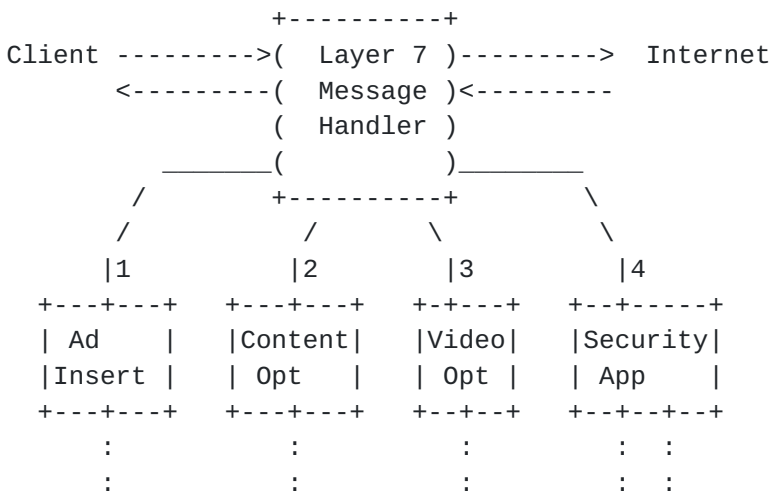

Figure 3: Layer 7 Service Chain Complexity

6. Conclusion and Recommendation

Service Chain touches upon Layer 2 to Layer 7. Challenges for Layer 4-7 service chain can be different from Layer 2-3.

This document provides common baseline for Layer 4-7 services and service chain and addresses their unique challenges.

7. Manageability Considerations

TBD.

8. Security Considerations

TBD.

9. IANA Considerations

This document requires no IANA actions. RFC Editor: Please remove this section before publication.

10. Acknowledgments

This draft has taken input from "Application Layer SDN" presentation given by John Giacomoni of F5 at Layer 123 conference. Thanks to Huang Shi Bi and Li Hong Yu for the valuable comments and suggestions.

This document was prepared using 2-Word-v2.0.template.dot.

11. References

[Application-SDN] J. Giacomonni, "Application Layer SDN", Layer 123 ONF Presentation, Singapore, June 2013

[SC-Use-Case] Liu, et, al., "Service Chaining Use Cases", < draft-liu-service-chaining-use-cases-00>, July, 2013

Authors' Addresses

Linda Dunbar
Huawei Technologies
1700 Alma Drive, Suite 500
Plano, TX 75075, USA
Phone: (972) 543 5849
Email: ldunbar@huawei.com

Donald Eastlake
Huawei Technologies
155 Beaver Street
Milford, MA 01757 USA
Phone: 1-508-333-2270
Email: d3e3e3@gmail.com

Intellectual Property Statement

Disclaimer of Liability

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.