Network Working Group J. You
Internet-Draft X. Wei
Intended status: Informational Huawei
Expires: December 30, 2015 June 28, 2015

Use Cases for SPUD
draft-yw-spud-use-cases-00

Abstract

The purpose of SPUD is to provide a standardized layer below the transport layer, to behavior as a communication channel between the host and network devices. So when a new transport protocol is deployed, it's easy for network devices to cooperate with it. New transports could have a common encapsulation to middleboxes. On the other hand, the transport layer could also make use of the state of network devices collected by the SPUD, to improve transport performance.

This document provides some exemplary use cases for SPUD, especially in stateful firewall and TCP optimization. The objective of this draft is not to cover all conceivable p2a or a2p signaling in detail. Rather, the intention is to explain the requirements in these use cases as far as it is required to complement the problem statement of the SPUD.

Requirements Language

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 [RFC2119].

Status of This Memo

This Internet-Draft will expire on December 30, 2015.

Table of Contents

1. Introduction

The purpose of SPUD is to provide a standardized layer below the transport layer, to behavior as a communication channel between the host and network devices. So when a new transport protocol is deployed, it's easy for network devices to cooperate with it. New transports could have a common encapsulation to middleboxes. For example, when a new transport protocol is deployed, the middlebox could be configured to allow the traffic even though it doesn't know the new transport protocol. On the other hand, the transport layer could also make use of the state of network devices collected by the

SPUD, to improve transport performance.

[I-D.hildebrand-spud-prototype] provides a prototype for grouping UDP [RFC768] packets together into a "tube". [I-D.hardie-spud-use-cases]

describes some basic use cases for path declarations (p2a) and application declarations (a2p).

This document provides some exemplary use cases for SPUD, especially in stateful firewall and TCP optimization. The objective of this draft is not to cover all conceivable p2a or a2p signaling in detail. Rather, the intention is to explain the requirements in these use cases as far as it is required to complement the problem statement of the SPUD.

2. Terminology

This section contains definitions of term frequently used throughout this document.

ICMP: Internet Control Message Protocol

MSS: Maximum Segment Size

RTT: Round-Trip Time

SPUD: Substrate Protocol for User Datagrams

TCP: Transmission Control Protocol

TCB: TCP Control Block

UDP: User Datagram Protocol

a2p: Application to Path

p2a: Path to Application

3. Use Cases

3.1. Firewall

Firewall is a kind of widely deployed middlebox that controls the incoming and outgoing network traffic based on an applied rule set; firewall usually works mainly on network layer and transport layer. Based on whether firewall keeps track of the state of network connection across it, firewall could be divided into stateless firewall and stateful firewall, the latter one is more commonly deployed. Compared with stateless firewall, stateful firewall maintains states for network connections, the state could include IP addresses, ports etc. Based on the states different packets could be associated with a session. Stateful firewall could provide more fine- grained network control, as the existing stateful firewall

could not only filter packet based on installed security policy but also based on state information. For the design of SPUD, firewall would be a very typical middlebox with which SPUD should be able to interact.

The process of dealing with packets by stateful firewall could be separated into two aspects: first, firewall administrator designs and installs a set of security policies on firewall, examples of policies could be prohibiting certain port numbers, prohibiting external host from starting connection to internal hosts protected by the firewall; second, if a connection is allowed, then firewall will establish a session for the connection, and the session might be updated as long as new packets belonging to the session arrive.

When a packet arrives at the firewall, if the packet belongs to an existing session in firewall the packet will be allowed to pass through; however if the packet doesn't belong to any existing session, then security policy rules will be applied to decide whether the packet is allowed, in case of the packet is allowed, a new session will be created, if not the packet will be discarded.

In order for firewall to maintain session state for network connection, the firewall must be able to know which session packets belong to, i.e. how to associate independent packets with a session. We should be aware that associating packets with a connection (we named it connection-A here) to one session (named session-A here) is just one case; another case is associating packet(s) not belonging to connection-A but related to connection-A with session-A, for example, normally firewall will forbid ICMP message (e.g. ICMP Host

unreachable or ICMP Network unreachable) initiated from external network from passing through the firewall, but there are cases that if the external initiated ICMP message is associated with an existing session, the ICMP message should be allowed; another example is association between FTP data connection and FTP control connection.

Another issue about firewall is that when a new transport protocol is designed, the legacy firewall will be unable to properly deal with traffic using the new transport protocol without any update of the firewall for the new transport protocol, and the situation usually leads to the traffic to be blocked. So one purpose of SPUD is to increase the deployment of new transport protocol even if the firewall is not updated for the new transport protocol. There could be three potential solutions for this: (a) define a fixed SPUD layer and hide the transport layer totally from firewall; (b) define a flexible SPUD layer that could provide transport protocol related information to firewall in a standardized form, so that the firewall could learn about the behavior of new transport protocol; (c) define

an extensible SPUD layer, and the new transport protocol will be designed by extending SPUD.

From the analysis above, in order to satisfy stateful firewall requirements, SPUD protocol needs to provide the following functions:

(1) Assist firewall to associate a set of packets to the same session, even though packets don't belong to but relate to the connection.

(2) Provide enough traffic related information for firewall to maintain state for the traffic.

(3) Indicate the start and stop of a session.

(4) Provide a standard interface to firewall assisting firewall to deal with new transport protocol.

3.2. TCP Optimization

TCP [RFC793] is a connection-oriented reliable transport protocol. Each TCP connection maintains state, usually in a data structure

called the TCP Control Block (TCB), containing information about the connection state, such as RTT, MSS, congestion avoidance threshold. These parameters can be shared across connections to the same host, since they are in fact per-path (per-host-pair) dependent [RFC2140].

The goal of state sharing is to improve transient transport performance. However, these parameters cannot clearly reflect the state of a specific on-path network device, for example, the state of network bottleneck device on the path, which is very useful information for TCP for TCB initialization. If TCP could obtain the available bandwidth of the bottleneck device, it could adjust the maximum congestion window size based on the ideal window size, which can be calculated by "the bottleneck bandwidth * RTT". [I-D.sprecher-mobile-tg-exposure-req-arch] presents the similar use case, i.e. the requirements for a mobile throughput guidance exposure mechanism that can be used to assist TCP in cellular networks, ensuring better performance.

SPUD can be used for path declarations: information delivered to the endpoints from devices along the path. Path declarations can be thought of as enhanced ICMP for transports using SPUD, allowing information about the condition or state of the path or the tube to be communicated directly to a sender. Therefore, this usage would enable the on-path network devices (e.g. gateway, router) to transmit their states (e.g. delay, packet loss rate, and bandwidth) to the transport layer. As the transport layer can obtain the state of the

on-path devices, it could make use of this information to configure initial transport parameters. The objective is to optimize the behavior of transport protocols. Meanwhile, the network state could be shared across connections if they share the paths. The possible scenario is shown in Figure 1.

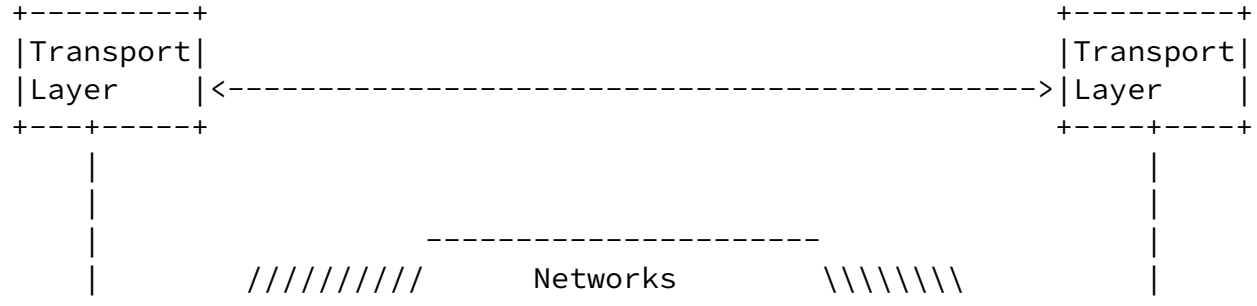

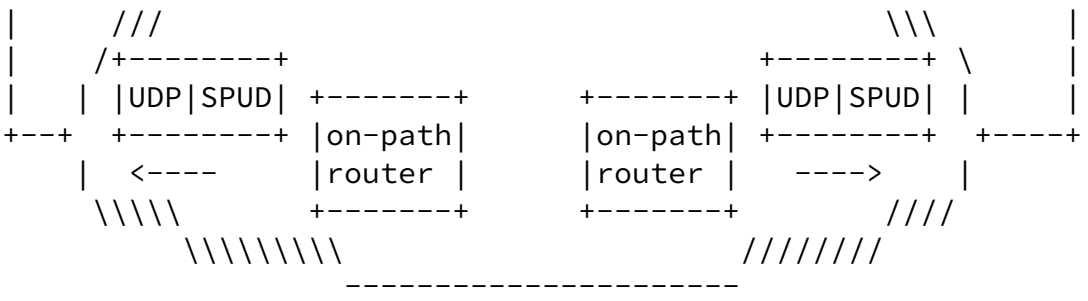

Figure 1: TCP Optimization Using SPUD

From the analysis above, in order to improve transport performance, SPUD protocol needs to provide the following functions:

(1) Provide the state (e.g. bandwidth, delay) of on-path network devices to the transport layer.

4. IANA Considerations

This document has no actions for IANA.

5. Security Considerations

For the use cases proposed in this document, the paritcipating entities need to have certain kind of trust relationships with each other. Meanwhile, the exposed information should be verifiable by each other. How to establish the trust relationship and how to verify the exposed information will be discussed in later versions of this document.

6. Acknowledgement

The authors would like to thank Mirja Kuehlewind and Michael Welzl for their comments.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2140] Touch, J., "TCP Control Block Interdependence", RFC 2140, April 1997.

[RFC3124] Balakrishnan, H. and S. Seshan, "The Congestion Manager", RFC 3124, June 2001.

[RFC768] Postel, J., "User Datagram Protocol", RFC 768, August 1980.

[RFC793] Postel, J., "Transmission Control Protocol", RFC 793, September 1981.

7.2. Informative References

[I-D.hardie-spud-use-cases]
Hardie, T., "Use Cases for SPUD", draft-hardie-spud-use-cases-01 (work in progress), February 2015.

[I-D.hildebrand-spud-prototype]
Hildebrand, J. and B. Trammell, "Substrate Protocol for User Datagrams (SPUD) Prototype", draft-hildebrand-spud-prototype-03 (work in progress), March 2015.

[I-D.sprecher-mobile-tg-exposure-req-arch]
Jain, A., Terzis, A., Sprecher, N., Swaminathan, S., Smith, K., and G. Klas, "Requirements and reference architecture for Mobile Throughput Guidance Exposure", draft-sprecher-mobile-tg-exposure-req-arch-01 (work in progress), February 2015.

Authors' Addresses

Jianjie You

Huawei
101 Software Avenue, Yuhua District
Nanjing 210012
China

Email: youjianjie@huawei.com

Xinpeng Wei
Huawei
No. 3, Xin-Xi Rd., Haidian District
Beijing 100095
China

Email: weixinpeng@huawei.com