RESEARCH DISCLOSURE
www.researchdisclosure.com

Certified Document

The Industry Standard Disclosure Publication Service

Method for In-band Meta-data Transfer

Disclosed anonymously

Research Disclosure database number 623051

Published in the March 2016 paper journal

Published digitally 18 February 2016 16:18 UT

Research Disclosure™ has been published monthly since 1960.
ISSN 0374-4353.

Research Disclosure, Questel Ireland Ltd
Killernan, Kilmaine, County Mayo, Ireland
Email: info@researchdisclosure.com
Phone: +353 (0) 93 33870

Method for In-band Meta-data Transfer

1 TECHNICAL AREA / KEYWORDS

Transport protocol, middlebox, explicit cooperation

2 BACKGROUND

2.1 Technical background / Existing technology

The transport protocols (TP) started evolving as a result of new requirements emerging from the evolution on the application and service layer. As a result, new TP designs have been proposed such as QUIC – a TP for HTTP2 [3].

There are a few requirements to TP evolution, which are summarized in a recent internet draft [4]. One is the possibility of communicating with the middleboxes. The reason is that middleboxes and the policies they enforce have an important role in the eco-system. The internet is a critical infrastructure a number of services are using from emergency services to banking systems, which should therefore be protected against misbehaving protocols. Middleboxes that filter out unknown TP implementations together with miss-behaving kernel-based TP implementations are the means to keep the Internet stable.

In conclusion, a solution is needed that lets the middleboxes not block “good” encrypted user space TP implementations but in a way that protects against misbehaving TPs. This is likely possible only via some explicit cooperation between the TPs and middleboxes. One proposal in IETF for such communication is a separate (TP independent) substrate protocol, SPUD ([1]-[2]).

Search for state-of-the-art patens:

Search words:

Thomson innovation search for (in-band signalling) OR (in-band signaling) AND TCP AND checksum:

US20140355623A1 (Futurewei): “Transmission Control Protocol (TCP) Connection Control Parameter In-Band Signaling”. A method for in-band signaling, the method comprising: receiving one or more packets over a transport connection control (TCP) connection extending between a source endpoint and a destination endpoint; embedding, by the destination endpoint, a sequence number pattern in one or more acknowledgment (ACK) messages, the sequence number pattern being pre-associated with an in-band signaling message; and sending, by the destination endpoint, the one or more ACK messages to the source endpoint.

US9049046B2 (Cisco): “System and method for offloading data in a communication system”. It makes use of GTP tunnel information for in-band tunnelling

Search for (middlebox AND checksum) gave no more relevant results

Search for ("invalid checksum” AND middlebox") gave the following relevant result below:

US8964755B2 (BT): “Obtaining information from data items”: the idea is for a sender to repeat a widely accepted protocol option that would not previously have been expected to be repeated. Such repetition can therefore act as a covert signal indicating to the receiver that one or both of the repeated fields should be interpreted with new semantics. Note that only the type of the protocol option is repeated, not necessarily all the parameters which may have different values. A receiver that understands the new protocol will be arranged to detect the repetition and interpret the modified semantics of the repeated-fields in the appropriate way. A receiver that has not been updated to understand the new protocol will read the first occurrence, act on it, then read the second occurrence and act on it, and so on, most probably without noticing it has processed a similar option already.

Search for ("invalid checksum” AND proxy"): gave no relevant results

Literature search for related proposals:

A. Yourtchenko (Cisco), “Introducing TCP Long Options by Invalid Checksum” IETF Internet Draft <draft-yourtchenko-tcp-loic-00>

- It involves use of a pair of connection initialization packets the first of which conforms to existing standards while the second uses an updated protocol in an extension to the header that would confuse a legacy receiver. This second initialization packet also uses a different checksum algorithm so that a legacy receiver will consider it invalid and discard it. The legacy receiver then treats the first packet as the only initialization packet. However, an updated receiver will understand the new checksum algorithm and the new protocol, so it will be able to process the second initialization packet. It will also understand the first packet

A. Jain et al. (Google-Nokia), “Mobile Throughput Guidance Inband Signaling Protocol” <draft-flinck-mobile-throughput-guidance-03.txt>

- Using the TCP options from a Middlebox to a Server to carry throughput guidance associated with an ongoing TCP connection

2.2 Problems with existing solutions

While SPUD is one potential solution for the problem, it has deployment and migration issues, because it requires support from both endpoints. A solution would be needed that works in the current eco-system.

An important observation is that meta-information should be conveyed at the Transport Protocol layer, because of the potential end-to-end encryption on the TP layer and above, and because of the threat of potential packet modification by the network below the TP layer (by NAT/FW). Potential problems with these solutions are:

- Preparing a signaling packet with the same identifiers (5-tuple) as a TP connection which the signaling refers to could cause problems to the TP implementations not aware of this method (e.g., connection break).

- An alternative would be preparing UDP signaling packets with the same IP address/port-pair: the main problem is that it cannot be ensured that it goes the same path as the data connection. Similarly, server load sharing may cause the signaling message arrive at a different server

3 BRIEF SUMMARY OF THE PROPOSED SOLUTION / ABSTRACT

We propose a method, where the signaling packets are sent in-band using the same connection identifiers (five-tuple), but prepared in such a way that the packet fails during the error check of the receiving connection endpoint. The methods to ensure that the error check fails depend on the way TPs perform error checking. The signaling packets would also have a 'magic number' easily identifying the new protocol.

The signaling messages are then fast discovered by the receiver by checking for the presence of a magic number, end-points not supporting this procedure simply discard these messages due to the failed error check.

4 ADVANTAGES OF THE PROPOSED SOLUTION

The advantage of the proposed method is that it makes in-band signaling to/from a Middlebox possible, with clear separation of middlebox signaling from e2e traffic through separate packets, which enables simple processing i.e., middleboxes do no need to re-assemble the packet (new checksum etc.) with data for the endpoint. Only the receiving end-point has to be upgraded for this procedure. Even if that end-point is not upgraded it does not cause protocol problems.

5 DETAILED DESCRIPTION OF THE PROPOSED SOLUTION AND FIGURES

The basic concept is depicted in the schematic Figure 1.**Erreur ! Source du renvoi introuvable.** In this example Middlebox 1 sends a signaling message to Server signaling logic, by using the identifiers of the connection between a Client and a Server, inferred in Step 2. Note that communication between any endpoint and any Middlebox may be achieved by the same method; for example, Middlebox 1 could send flow-related meta-information to another Middlebox 2 rather than the Server signaling logic. Two examples of how to prepare the signaling packets based on the identified TP in Step 3 are given in sections 5.1 and 5.2, respectively.

5.1 Preparing signaling packets for a TCP connection

Note that the method described here may be applied for other TPs that have an open CRC too.

The idea is that the sender of the signaling message prepares a "checksum" field in the TCP header that will fail the CRC check of the connection endpoint. This is done by e.g., XOR-ing the would-be CRC with a magic number. Thus, legacy servers will drop the packets (since there is very low probability of false positives). The wire format if the end-to-end data vs. the signaling packets is shown in Figure 2.

Note that it is advisable that the meta-data starts with the so-called "magic number". The magic number is a fixed number that is sufficiently large to reduce the likelihood of having the same pattern in the corresponding part of the end-to-end data packet [Note: if it replaces textual parts in the packets e.g., non-encrypted TCP, it may be chosen such that to be invalid UTF-8, UTF-16 (both big- and little-endian), and UTF-32 (both big- and little-endian) thus avoiding accidental collision completely]. The purpose of the magic number is to allow a very fast (i.e., trivially implementable in hardware) way to decide that packets that do not include the magic number are NOT signaling packets.

5.2 Preparing signaling packets for a QUIC connection

Note that the method described here may be applied for other TPs that have a CRC in the encrypted part of the TP packets.

The solution builds on the assumption that, by manipulating the packet, the byte string in the place of the encrypted CRC will not be a valid CRC, which will show that this packet does not belong to the normal e-2-e communication and be therefore dropped by a legacy server. The sender of the signaling packet may thus just re-use a packet passing through in the past and replace the encrypted part with meta-data starting with a magic number identifying the protocol. This is shown in Figure 3. In the example given in the figure, there is a CRC/authentication after the meta-data that is used for integrity protection/authentication of meta-data.

Note that it is practical to use sequence numbers in the meta-data communication that have been used already, which further secures that the other endpoint will drop the signaling packet if receiving it.

5.3 Figures

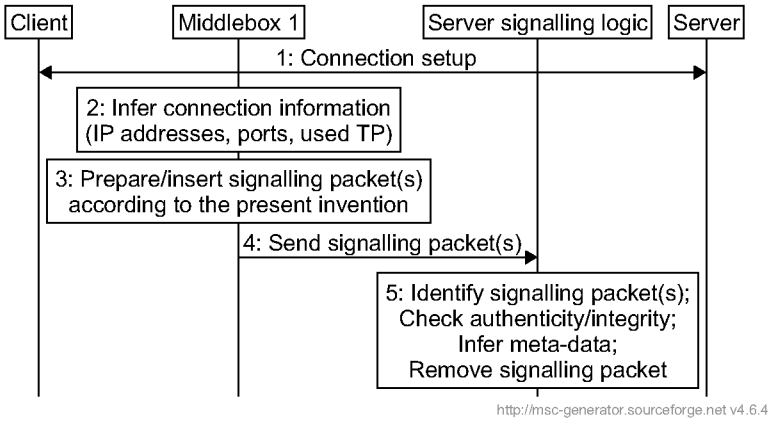

Figure 1 Example sequence diagram illustrating the proposed in-band meta-data communication

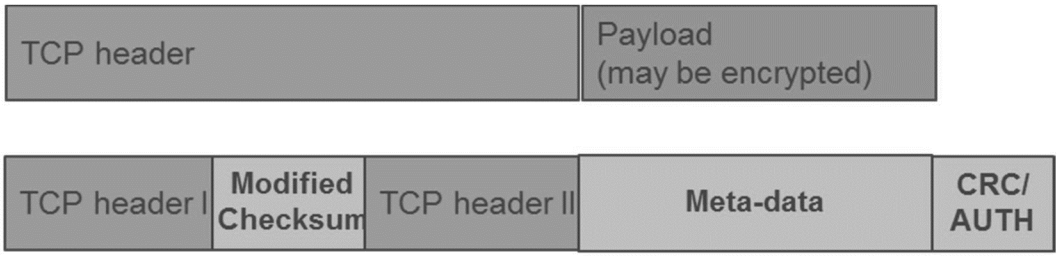

Figure 2 Wire format of an end-to-end TCP data packet (top) and a signaling packet for the same TCP connection (bottom) by using the proposed method.

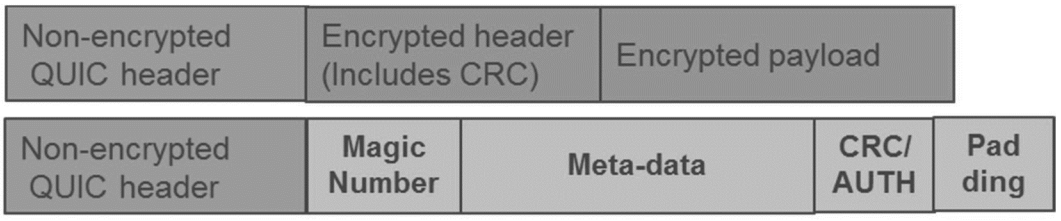

Figure 3 Wire format of an end-to-end TCP data packet (top) and a signaling packet for the same TCP connection (bottom) by using the proposed method

6 CORE ESSENCE OF THE SOLUTION

A method for in-band meta-data communication, where the signaling packets use the same connection identifiers (five-tuple) as the end-to-end TP connection, but are prepared in such a way to fail the error check of the receiving TP endpoint

Where the preparation of the signaling packets is based on the identified TP characteristics

- If the TP has an encrypted CRC then a past packet is reused and the encrypted part is replaced by the signaling meta-data
- If the TP has an open CRC then a new CRC is prepared that will fail the error check of the receiving endpoint with high probability
    - Where the new CRC is calculated as the XOR of the correct CRC and a magic number

7 ABBREVIATIONS

Abbreviation	Explanation
CRC	Cyclic redundancy Chech
FW	Firewall
NAT	Network Address Translation
SDU	Service Data Unit
SPUD	Substrate Protocol for User Datagrams
TCP	Transmission Control Protocol
TP	Transport Protocol
UDP	User Datagram Protocol

8 REFERENCES

Reference	
[1] SPUD mailing list, see http://www.ietf.org/mail-archive/web/spud/current/maillist.html	
[2] Brian Trammel, Substrate Protocol for User Datagrams https://www.ietf.org/proceedings/92/slides/slides-92-spud-1.pdf	
[3] QUIC: A UDP-Based Secure and Reliable Transport for HTTP/2, http://tools.ietf.org/html/draft-tsvwg-quic-protocol-00	
[4] Enablers for Transport Protocol Evolution, presentation at IETF93, https%3A%2F%2Fwww.ietf.org%2Fproceedings%2F93%2Fslides%2Fslides-93-tsvarea-0.pdf	

Disclose Anonymously