DOTS
Internet-Draft
Intended status: Standards Track
Expires: August 6, 2016

R. Moskowitz
J. Xia
Huawei

February 3, 2016

DOTS over GRE
draft-moskowitz-dots-gre-00.txt

Abstract

This document describes using a GRE tunnel to deliver DOTS messages between DOTS agents and compares it to other methods.

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 6, 2016.

Table of Contents

1. Introduction

This document describes using a GRE [RFC2784] tunnel to deliver DOTS messages between DOTS agents. Various alternatives for transporting DOTS messages are analyzed and the justification of GRE over alternatives as UDP over IP and UDP over ESP over IP is presented.

The intent of this document is to encourage discussion on the most effective set of protocols to provide the high reliability requirement spelled out in the DOTS requirements document [I-D.ietf-dots-requirements].

2. Terms and Definitions

2.1. Requirements Terminology

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 [RFC2119].

3. Problem Space

3.1. The Network issues faced by DOTS and UDP

DOTS messaging needs to occur during the worst time to expect reliable packet delivery. That is during a DDoS attack. Not only is link to (and potentially from) the DOTS client fully congested with the attack, but also the ISPs between the attacked DOTS client and the responsible DOTS server may have instituted UDP blocking mitigation activities.

It is this UDP blocking mitigation action that presents a double-edged effect. It lessens the impact of the attack, allowing TCP-based activity to continue. It stops any attack management, i.e. DOTS, messaging over UDP to traverse the portion of the network where the blocking is in affect.

3.2. Peer-to-peer not RESTful

A second problem, or more a challenge, in DOTS messaging is that it is really a peer communication. That is the DOTS server may be messaging the DOTS client at any time, including during an attack. Thus a client-service approach like RESTful would require 2 uni-directional sessions.

One example of a DOTS server message is an "Attack seems over" message from the server to the client.

3.3. Security Context

3.3.1. Stateful Security Context

Security Context is the collection of information to manage the securing of information. In this case DOTS messages. The only viable method for stateless security is secure data objects as in PEM [RFC1421]. Stateless security is very resource intensive and typically avoided unless it is the only effective approach. DOTS messaging will use a secure data channel which is stateful. This state needs to be managed and protected.

3.3.2. Security Context and Fate Sharing

Security Context often contains communication protocol information like IP addresses and transport ports. In these situations the security context is said to "share fate" with these aspects of the communications. If something disrupts the communication state, it disrupts the security context, often requiring some degree of security re-initialization.

The greater the fate sharing, the more rigid the security context and more prone to attack. Thus a secure message transport design goal is to lessen the degree of fate sharing.

4. Protocol Selection Considerations

Based on Section 3, DOTS messaging should take advantage of protocols that:

- Are bi-directional

  One such protocol is often used for bi-directional messaging is TCP. This is not a viable option as the ACK from sending a message from the DOTS client to server over the potentially uncongested uplink may never get back to the client over the congested down link.

- Are Not Commonly blocked, particularly during a DDoS attack

  UDP and ICMP fall into this avoidance category.

- Have minimal overhead

  DOTS messages that are sent during an attack should fit into a single MTU. The lower the protocol byte overhead, the more space available for the DOTS message itself.

- Are only enabled by need on a system

  It would be advantageous that the DOTS communication uses a protocol that is typically quickly discarded by most targeted systems. Even though these protocols are used by the DOTS agents, the DOTS agents will be hard to find to attack and will tend to have more resources available to deflect direct attacks.

- Support peer communications

  At all protocol levels, there must be no complexities in implementing peer communications. Pairing two uni-directional protocols to achieve this should be avoided.

  The security context that protects the DOTS messaging must support peer communications. That is a single DOTS agent security agreement would provide the complete context for DOTS security. Examples include IKEv2 [RFC5996] and HIPv2 [RFC7401]. It is noted that these maintain two uni-directional Security Associations within the security context to properly manage the key usage in each direction.

o Provide secure communications with minimal fate-sharing

The security context should be resilient to DOTS agent restart and thus potential loss of protocol state. At best there should be no fate-sharing with any protocol state. An option for security state to be stored in a safe manner so that it need not be renegotiated after agent restart makes forcing an agent restart an uninteresting attack.

5. The DOTS Protocol Stack

Below are three possible protocol designs. The compressed GRE design, Section 5.2, best meets the selection considerations (Section 4).

5.1. GRE full stack tunnel

GRE is basically used to tunnel Ethernet payloads across an IP network. For example an IPv4 datagram can be tunneled within GRE with a GRE Protocol Type of 0x800. This is simple to implement on a system, as GRE appears to IP as an interface. DOTS messaging can secured with SSE [I-D.moskowitz-sse] on UDP over IPv4 or IPv6 within this GRE tunnel.

GRE can also work well in a NAT traversal deployment scenario.

5.1.1. Design Analysis

The per-packet byte cost of GRE and an inner IP envelope (IPv4 or IPv6) is balanced in part by the envelope simplicity of SSE. SSE has the advantage of being completely free of fate-sharing with the lower protocol levels. GRE, as indicated, is relatively easy to support as a psuedo-interface. This is weighed against SSE being new, and any Key Management Protocol would need negotiation parameters to support SSE.

Use of SSE also allows secure transport of DOTS messages over non-IP connections, for example SMS. The low SSE envelope overhead of as little as 20 bytes can allow for 120 bytes for a single SMS message. SMS message continuation can allow for longer DOTS messages.

5.2. GRE with compressed stack tunnel

The full GRE stack approach may overly constrain the size of the DOTS message that can fit within a single MTU. There are approaches to compress this into a smaller size.

There are two approaches to reduce the header overhead of the GRE full stack tunnel outlined above. RObust Header Compression [RFC3095] is the well-known approach. Within this compression, the datagram will logically be the same as above.

The actual inner IP header could be compressed to zero bytes by using the same source and destination addresses of the outer IP header. This is more than specified in ROHC, and would involve additional specification. NAT traversal design considerations need to be included in the compression scheme.

5.3. ESP transport mode

ESP [RFC4303] in transport mode (or BEET with HIPv2) Provides a familiar approach to protect UDP traffic. ESP with IKEv2 fate-shares with both IP and UDP. ESP with HIPv2 only with UDP. Either way, loss of UDP state due to a DOTS server crash would require reestablishment of the security state. This keeps attacks against the DOTS server as an important attack surface to weigh against the familiarity of ESP with IKEv2 or HIP.

ESP limits secure DOTS messaging to IP networks. A different method would be needed for sending DOTS messages over SMS or require IP over a modem connection.

ESP NAT traversal uses UDP and thus reintroduces the UDP blocking concern discussed above.

6. Management Considerations

6.1. DOTS agent connectivity management

A DOTS client needs to be configured with knowledge of the DOTS servers. This may either by an IP address or an FQDN. If FQDN is used, IP addresses should be cached as DNS lookups may fail during an attack.

6.2. Secure Context management

Some trustworthy authentication needs to be set up on both sides. This authentication knowledge will be used by a Key Management Protocol like IKEv2 or HIPv2 to create the security context. Either can manage the security context for ESP or SSE. Two strong authentication methods use digital certificates or raw public keys.

Digital certificate trustworthiness may not be easy to determine. There are many issues such as which Certificate Authority to trust and how to manage Certificate Domain trust leakage. These issues

often result in needing to manage an authorization list of trusted certificates.

Raw public keys for IKEv2 [I-D.kivinen-ipsecme-oob-pubkey] or HIPv2 HITs can be managed in an ACL without the cost associated with Digital Certificates. Replacing 'old' keys can be associated with the DOTS business model of contract renewal.

7. IANA Considerations

No IANA considerations exist for this document at this time.

8. Security Considerations

A DDoS attacker would greatly benefit from disabling DOTS. This may be accomplished by:

- Blocking DOTS traffic.
- Disabling DOTS servers.
- Disabling DOTS clients.

A key component of this proposal is to lessen the likelihood of ISPs from blocking DOTS traffic by not using UDP. Whatever protocol DOTS uses, may be used in future DDoS attacks, but will not be as effective as UDP based attacks. Thus not using UDP is a worthwhile goal.

DOTS server resiliency to attacks is a critical goal. Loss of a DOTS server can impact many clients (customers). The less fate-sharing the higher the attack resiliency, which is why this document recommends the GRE with compressed stack tunnel, Section 5.2, approach.

DOTS clients will tend to be invisible to attackers, but over time they will be discovered for targeted attacks, thus the same resiliency considerations applied to the servers also apply to the clients. Additionally, DOTS clients should avoid access to as many Internet services as possible, as at critical times they may be blocked. Thus a non-PKI authentication scheme as in raw public keys has the advantage of needing one less Internet resource that may be blocked.

9. Contributors

The following contributed actively to the this document: Sue Hares (Huawei)

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

10.2. Informative References

[I-D.ietf-dots-requirements] Mortensen, A., Moskowitz, R., and T. Reddy, "DDoS Open Threat Signaling Requirements", draft-ietf-dots-requirements-00 (work in progress), October 2015.

[I-D.kivinen-ipsecme-oob-pubkey] Kivinen, T., Wouters, P., and H. Tschofenig, "Generic Raw Public Key Support for IKEv2", draft-kivinen-ipsecme-oob-pubkey-14 (work in progress), October 2015.

[I-D.moskowitz-sse] Moskowitz, R., Faynberg, I., <>, H., Hares, S., and P. Giacomin, "Session Security Envelope", draft-moskowitz-sse-01 (work in progress), January 2016.

[RFC1421] Linn, J., "Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures", RFC 1421, DOI 10.17487/RFC1421, February 1993, <http://www.rfc-editor.org/info/rfc1421>.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <http://www.rfc-editor.org/info/rfc2784>.

[RFC3095] Bormann, C., Burmeister, C., Degermark, M., Fukushima, H., Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le, K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K., Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed", RFC 3095, DOI 10.17487/RFC3095, July 2001, <http://www.rfc-editor.org/info/rfc3095>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <http://www.rfc-editor.org/info/rfc4303>.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, DOI 10.17487/RFC5996, September 2010, <http://www.rfc-editor.org/info/rfc5996>.

[RFC7401] Moskowitz, R., Ed., Heer, T., Jokela, P., and T. Henderson, "Host Identity Protocol Version 2 (HIPv2)", RFC 7401, DOI 10.17487/RFC7401, April 2015, <http://www.rfc-editor.org/info/rfc7401>.

Authors' Addresses

Robert Moskowitz
Huawei
Oak Park, MI 48237
USA

Phone: +1-248-968-9809
Email: rgm@htt-consult.com

Jinwei Xia
Huawei
101 Software Avenue
Nanjing, Yuhua District 210012
China

Phone: +86-025-84565890
Email: xiajinwei@huawei.com