Messaging, Malware and Mobile Anti-Abuse Working Group
M³AAWG Introduction to Traffic Analysis

June 2016

Introduction

Protecting against pervasive monitoring and the use of encryption continues to be a major focus for the messaging industry. M³AAWG has already published initial recommendations for deploying TLS,[1] mitigating Man-in-the Middle attacks[2], and using forward secrecy to secure data[3] to help the messaging community understand how to better secure email in transit. Now M³AAWG would like to bring awareness to a different type of risk – a form of attack called *traffic analysis*. In this paper, we outline the key characteristics of traffic analysis, discuss potential ways to avoid it, and consider the advantages and disadvantages of deploying preventative measures.

Understanding Traffic Analysis with Respect to Messaging and Network Traffic

The content of messages encrypted with PGP/GPG (GNU Privacy Guard)[4] or S/MIME[5] is generally highly resistant to eavesdropping. Even if a third party manages to get a copy of a PGP/GPG-encrypted email (or an S/MIME-encrypted email), they are not likely to be able to decrypt and read it. However, even messages that are perfectly protected with end-to-end encryption remain potentially subject to traffic analysis attacks.

To understand the difference, consider the following summary table of email message elements visible to an intermediary SMTP server utilizing TLS for transmitting messages and their availability for traffic analysis purposes:

Email message elements	Vulnerable to traffic analysis?
Return-Path: header	Yes
Received: headers	Yes
From: header	Yes
To: header	Yes
CC: header	Yes
Date: header	Yes
Subject: header	Yes
Message-ID: header	Yes
Any/all other headers	Yes
Size of the message	Yes
Time message was received	Yes
Apparent encryption used by message	Yes
Message contents (assumed to be possibly or actually encrypted)	No

In a traffic analysis attack, the focus is not on the content, but on the message headers and other externally-observable artifacts associated with the message or the communication process itself. The summary table

above lists the elements of email messages that are visible to an intermediary SMTP server using TLS for transmitting messages. These are vulnerable to traffic analysis.

Similarly, traffic analysis of network traffic in general (rather than email traffic in particular) normally relies on the data usually provided by NetFlow[6] or equivalent protocols, as seen in the table below. This flow, too, is vulnerable to traffic analysis.

Network traffic flow	Vulnerable to traffic analysis?
# of packets in the flow	Yes
# of octets in the flow	Yes
Start/stop timestamps	Yes
Src/dst IP addresses	Yes
Src/dst port numbers	Yes
IP protocol	Yes
Type of Service value	Yes
Flags	Yes
[etc]	Yes
Packet contents	No

How Traffic Analysis Can Work

While it might seem as if a traffic analyst is heavily handicapped by not having access to a message's content, a surprising amount of information can be gleaned from non-content-related characteristics.

Let us consider a few examples of how traffic analysis can work:

1. ***Source/Destination Analysis:*** Imagine that a government employee working in a sensitive position sends an encrypted message to an investigative journalist or to a representative of a hostile foreign intelligence service. If that message gets noticed, the sheer fact that the employee has sent *any* such message (regardless of what the encrypted message might actually contain) would likely be enough to trigger a review. All that is needed to flag the communication for closer scrutiny is knowledge of the source and destination of the message.

   Or imagine an enterprise administrative system that processes credit card transactions. It normally talks only to an order-taking website and to the credit card company. If an analyst suddenly notices an encrypted connection between that computer and a heretofore unknown system elsewhere in the world, it would normally be a presumptive sign that something bad had happened and a sensitive system may have been compromised.

   A third example relates to botnet command and control hosts. If a desktop system suddenly begins communicating with a known botnet command and control host, that communication is a strong indicator that the desktop system may be botted. (This process is sometimes referred to as using a command and control host as a "finger-pointing tool.")

   So-called "pen register" and "trap-and-trace"[7] orders are a fourth example of source/destination traffic analysis. Pen register and trap-and-trace orders allow a judicially-authorized investigator to get information about the *source and destination* of specified calls, but

not the *content* of those calls. Pen register and trap-and-trace orders have routinely been used by law enforcement officers to obtain the telephone numbers of associates called by known members of organized crime or to identify distributors working with a known drug kingpin.

2. ***Traffic Volumetric Analysis:*** Investigators may not even need to know who is talking to whom – they just need to know how much talking they are doing. Simple traffic volumetric changes may convey important information in and of itself.

   Consider analysts monitoring encrypted military communications to and from a hostile country. Over a period of months or years, analysts tend to become familiar with routine baseline traffic: perhaps a typical small site exchanges 400 to 500 messages a day. Analysts come to know this is normal for that site.

   However, if that country suddenly begins to prepare to go to war, traffic levels will tend to *increase*, at least for the units that are being mobilized, as instructions are sent directing them to prepare for combat. Perhaps selected bases now begin to send or receive four or five times normal traffic levels. This simple volumetric change can be easily detected by an external monitor, even if the contents of the messages themselves are encrypted and cannot be deciphered. (This is the "increased chatter" phenomenon sometimes mentioned in media reports.) Because some bases have increased traffic levels while others do not, not only do analysts know that something is being planned, they also know where something is being planned and which units are likely being mobilized. This is valuable intelligence, even if they cannot see the actual message traffic.

   Alternatively, if normal traffic levels suddenly drop to nothing – i.e., "radio silence" is imposed in an effort to avoid potentially compromising unencrypted transmissions – that *downward* change in volume can also be an important indicator that something is imminent.

3. ***Sequence Analysis:*** Assume a site operates an encrypted mailing list server. External analysts monitoring network traffic to that server might see an encrypted message come into that server, and then shortly thereafter, they might see re-encrypted messages of roughly the same size go out to three dozen people at a variety of sites. This pattern is repeatedly observed.

   Given this pattern, the analysts can reasonably infer that the three dozen people are all part of the same encrypted mailing list and share an ongoing interest in some particular topic. For example, if a large fraction of the identified members are known to be involved with a political opposition group, a reasonable inference might be that *all* members of that list are associated with the opposition group, even if some of the group members were heretofore unknown as such.

4. ***Inferential Analysis:*** An analysis of possibly independent events, where the events constitute the "traffic," can create a better picture of some overarching event. Consider a spouse whose fortieth birthday is approaching and who discovers an email from a tent rental company, a receipt from a disc jockey and an entry on the telephone caller ID from a local bakery. These events do not indicate a direct act but putting together these bits of information can reveal the possibility of a surprise birthday party. This scrutiny and other methods make up the process of traffic analysis, allowing analysis of communications to occur even when the content of those communications is completely encrypted.

Avoiding Traffic Analysis Attacks

There are many ways to avoid traffic analysis attacks. We consider a few here to give a sense of the countermeasures that some have employed.

1. ***Decoupling Source and Destination via Intervening Hops:*** This countermeasure strives to avoid directly tying the sender and receiver of a communication by indirectly routing communications between them via multiple additional intervening hops. This is the approach that anonymous remailers generally employ. A good discussion of anonymous remailers[8] can be found in Wikipedia.

2. ***Tunneling All Traffic to a Third Party:*** Another approach to discourage traffic analysis is to tunnel all traffic from its actual source to a third party exit node, perhaps via a commercial VPN (Virtual Private Network) connection. The VPN provider must be selected with caution, as the provider will typically know both the identity of the VPN purchaser and, potentially, the contents of the traffic being tunneled over the VPN service. Furthermore, a network traffic analyst can easily detect such a countermeasure and authorities might attempt to block it by administrative or technical means. From a technical point of view, however, a local traffic analyst will see only a tunnel passing encrypted traffic to a single remote destination, while a remote traffic analyst will see only traffic going to a tunnel endpoint, not to its actual ultimate destination.

3. ***Decoupling Source and Destination via Broadcast or Publication Methods:*** This countermeasure strives to avoid tying the sender to the receiver of a communication by eliminating the direct connection between them, substituting broad publication as a replacement for targeted delivery of a message. As a simple non-online example of this, assume sender A wants to communicate with receiver B. Sender A could just send receiver B a letter but that would directly link the two, at least if the letter happens to gets noticed. As an alternative, A might instead place the communication to B in the form of a coded newspaper want ad.

   Since A is communicating publicly via a broadcast mechanism – the newspaper – it becomes impossible for a third party to know which of potentially hundreds of thousands of newspaper subscribers A is actually trying to communicate with. The link between A and B thus becomes difficult or impossible to establish.

   The online analog of publishing a message in the newspaper can be seen in Usenet News. Usenet News is a decentralized discussion platform organized around tens of thousands of different newsgroups (topical discussion areas). There are discussion groups for talking about gardening and other hobbies, discussion groups for technical topics such as programming languages, discussion groups for politics and religion and much more.

   Posts made by participants in a particular newsgroup are automatically "flooded" or "pushed" via news feeds between news servers carrying a particular group. This means that a message posted on the west coast of the United States can propagate to thousands of news servers all around the world in just a matter of seconds.

   Once a given post has been received by a news server, only the administrator of each of the thousands of local news servers can tell which user is reading a particular newsgroup from their server, if indeed anyone is reading a particular group at that site at all. Clearly it is

wasteful to broadcast an encrypted personal message worldwide when it is meant for only one person and only one person can read it, but this happens routinely nonetheless.

See, for example, the following excerpt from a posting made to the Usenet News newsgroup alt.anonymous.messages:

```
Path: ks2ni12827igb.0!nntp.google.com!peer01.iad.highwinds-media.com!
          news.highwinds-media.com!feed-me.highwinds-media.com!
          usenet.blueworldhosting.com!feeder01.blueworldhosting.com!
          news.mixmin.net!.POSTED!not-for-mail
From: "Nobody" <nor...@mixnym.net>
Newsgroups: alt.anonymous.messages
Subject: 25063f0266e3ab0a7f3d7935ee6d0c89acdca564b3d32c17
Date: Wed, 5 Nov 2014 14:04:03 -0600
Organization: Mixmin
Lines: 52
Message-ID: <m3dvrj$ed5$1@news.mixmin.net>
Injection-Date: Wed, 5 Nov 2014 20:04:03 +0000 (UTC)
Injection-Info: news.mixmin.net;
          posting-host="8Fb10CoLdDqohFEVzZq91qcbCxo";
          logging-data="14757"; mail-complaints-to="ab...@mixmin.net"
X-Received-Bytes: 3749
X-Received-Body-CRC: 1349451157

-----BEGIN PGP MESSAGE-----

hQIMAwAAAAAAAAAARAAnuPu5p0HtqDtaxrYHBkH1q684HfdOr+OfNEH8DjG6XE6
[continues]
```

Anyone can see the message, but because it is PGP-encrypted, only the person with the appropriate private key can decrypt and read it. Moreover, because it is sent worldwide as part of a newsgroup, rather than being sent directly from the author to the receiver, no one knows who originated the message or who is meant to receive it. (While the message is PGP-encoded, it is not encrypted with a normal PGP key closely tied to an actual identity.) This approach is quite strong but it is not perfect. See, for example, the discussion in "De-anonymizing Alt.Anonymous.Messages."[9]

4. ***Constant Volume Traffic Channels:*** To avoid volumetric traffic analysis attacks, stations might arrange to exchange a constant volume of encrypted traffic, sending readily-ignorable, random content when not transmitting real messages. Because the volume of traffic is invariant, volumetric approaches are frustrated. Constant volume traffic channels also have the advantage of proving that a given communication path is live and usable at all times.

   The disadvantage of this approach is that communication bandwidth ends up being permanently committed, thereby wasting capacity. This approach also tends to make it hard not to notice an ongoing, presumably important, connection between the two sites. Finally, sending a constant stream of traffic increases the corpus of encrypted text available for potential attacks by cryptanalysts.

   Alternatively, one could imagine simulated traffic bursts mimicking what might be seen if an actual event of interest were taking place. If there were enough of these false alarms, a real uptick in traffic volume might easily be mistaken for another simulated burst of traffic.

Consequences of Attempts to Evade Traffic Analysis

Measures taken to avoid traffic analysis attacks tend to make a particular provider's traffic stand out from routine network traffic. Traffic that might have been routinely ignored will now more likely be subject to scrutiny.

Because traffic sent through anti-traffic analysis channels such as anonymous remailers cannot be attributed, spammers may attempt to abuse such sites to deliver their messages. Operators of anonymous remailers generally resist abuse of their services, but their ability to police misuse is imperfect at best. As a result, some receiving systems may routinely block all traffic from known anonymous remailers as a matter of policy.

Anonymous remailers may intentionally introduce delays in retransmitting messages, seeking to decouple network-observable input traffic from network-observable output traffic. This additional latency is not a problem normally but it can be in some time-critical situations.

Many anti-traffic analysis solutions rely on sending traffic through intervening nodes. If unencrypted messages are rerouted via one or more untrustworthy intervening nodes, those nodes may eavesdrop on that traffic. Untrustworthy intervening nodes can also confound traffic delivery by dropping traffic that they have been entrusted to retransmit.

Counter-traffic analysis measures may result in traffic being routed overseas. This may mean that:

- Domestic traffic that is normally exempt from analysis due to statutory exclusions may become subject to review since the domestic origin of the traffic has been obfuscated and the traffic appears to be international in origin.

- Traffic may become subject to the remote site's legal and regulatory regime as well.

- Search engines and other resources may inappropriately localize the contents based on the endpoint of the tunnel being used; e.g., if a VPN terminates in Germany, the end user may be shown content in German rather than English.

- Indirectly routing traffic will also introduce unavoidable actual latency simply because of the distance that the traffic must travel to/from a remote tunnel endpoint.

- Commercial solutions that are meant to inhibit traffic analysis (e.g., commercial VPN services) may be expensive to purchase at scale.

Conclusion

Most users have limited awareness of traffic analysis as a risk and may be disinclined to employ measures meant to counter traffic analysis. Some countermeasures may also unintentionally increase scrutiny applied to traffic. Making traffic more resistant to analysis can be challenging. It requires effort from the groups that design our standards at network, protocol and operational levels, as well as user awareness. These longer-term goals can begin today, with experimentation and deployment as they are proved to be reliable and trustworthy. Traffic analysis may be a more specialized problem, but M^{3}AAWG would like to bring awareness to this risk/attack scenario and to recommend that the messaging industry continue working on countermeasures to mitigate it.

References

[1] TLS for Mail: M³AAWG Initial Recommendations, Messaging, Malware and Mobile Anti-Abuse Working Group, December 2014, https://www.m3aawg.org/sites/default/files/document/M3AAWG_TLS_Initial_Recommendations-2014-12.pdf

[2] M³AAWG Initial Recommendations for Addressing a Potential Man-in-the-Middle Threat, Messaging, Malware and Mobile Anti-Abuse Working Group, July 2015, https://www.m3aawg.org/sites/default/files/M3AAWG-Man-in-the-Middle-Recommendations2015-07.pdf

[3] M³AAWG Initial Recommendations for Using Forward Secrecy to Secure Data, Messaging, Malware and Mobile Anti-Abuse Working Group, January 2016, https://www.m3aawg.org/sites/default/files/m3aawg-forward-secrecy-recommendations-2016-01.pdf

[4] "The Gnu Privacy Guard," GnuPG, https://www.gnupg.org/

[5] "Client Certs and S/MIME Signing and Encryption: An Introduction," Joe St Sauver, February 20, 2012, https://www.stsauver.com/~joe/maawg24/maawg24.pdf

[6] "NetFlow," http://en.wikipedia.org/wiki/NetFlow

[7] "Trap and Trace Device," https://en.wikipedia.org/wiki/Trap_and_trace_device

[8] "Anonymous Remailer," http://en.wikipedia.org/wiki/Anonymous_remailer

[9] "De-anonymizing Alt.Anonymous.Messages," Tom Ritter, Ritter.vg, August 3, 2013, https://ritter.vg/blog-deanonymizing_amm.html

As with all best practices that we publish, please check the M³AAWG website (www.m3aawg.org) for updates to this document.

M3AAWG104