Vahab Pournaghshband
University of San Francisco
vahab.p@usfca.edu

LECTURE
1

On Prevalence of Translucent Middleboxes

1.1 What Are Middleboxes?

Abstractly, the Internet follows the end-to-end principle, with smart endpoints and a dumb network. However, with the emergence and rapidly growing prevalence of middleboxes deployed at various points in the network, the actual Internet is far more complex.

A middlebox, defined as "any intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host [6]," manipulates traffic for purposes other than simple packet forwarding. In addition to routing the traffic, middleboxes are potentially capable of making serious changes to what happens to a traffic flow on the network. These changes have meaningful implications to the senders and receivers.

A wide variety of middleboxes have been proposed, implemented and deployed during the last decade [18]. Today's enterprise networks rely on a wide spectrum of specialized applications of middleboxes. Middleboxes come in many forms such as proxies, firewalls, IDS, WAN optimizers, NATs, and application gateways and for various purposes including performance and security improvement and compliance. They are the integral part of today's Internet and play an important role in providing high levels of service for many applications. Recent papers have shed light on the deployment of these middleboxes [22, 18] to show their prevalence. And a recent study [19] shows that the number of different middleboxes in an enterprise network often exceeds the number of routers. Trends such as proliferation of smartphones and wireless video are set to further expand the range of middlebox applications [1].

1.2 Why is Detecting Middleboxes Important?

Knowing the existence of such influence could be beneficial and in some cases crucial to the end-hosts. Sometimes the end-hosts would behave differently based on what they sense about what is happening to their traffic. In such cases, an accurate detection of what is happening to their traffic is the first step. Here, to illustrate this idea, we present a few

scenarios from different categories.

Scenario I. Performance

The sender is about to send highly compressible data, making the compression worthwhile. The sender checks if an end-to-end compression or link compression on bottlenecks on the path is deployed. If they detect that compression is already in place, they would not compress the traffic stream, as double compression is redundant and costly. In a similar scenario, the sender might not encrypt if they detect that strong encryption is already in place by the gateways.

Scenario II. Security

The sender knows the receiver is using a wireless connection, but is unsure if that connection is secure. If he detects that the last link is unencrypted, he would either refrain from sending sensitive information or would apply end-to-end encryption to the channel. For example, Facebook, until recently, by default, did not provide end-to-end SSL encryption to its users, perhaps due to lack of available resources required to encrypt all users' contents for all users. Facebook servers, to use their resources effectively while protecting Facebook users' privacy, could first sense whether the user is using a secure wireless connection or not, and then apply end-to-end encryption only if the user needs that protection. Conversely, the receiver would mark the incoming data as untrusted if he detects the sender's wireless link is insecure.

Detecting the presence of a Man-in-the-Middle (MITM) is the first step to take any action against it or apply the necessary protection to their network flow to make potential attacks by the attacker ineffective. Therefore, it is crucial for the end-hosts to be able to detect the MITM's presence.

Scenario III. Politics

An Internet user in an oppressive country might detect Internet censorship imposed by their ISP and then chooses to use a proxy to bypass it. In a different scenario, an Internet user detects wire-tapping on their network and uses evading techniques such as Tor or VPN to make wire-tapping less effective.

Scenario IV. Debugging

Despite their growing importance in handling operational traffic, middleboxes are notoriously difficult and complex to manage [18]. Hence, they are more prone to become a point of failure independent of network conditions. It is worth understanding if a problem with the middlebox (e.g., malfunction, misconfiguration, or overload) is causing the network to misbehave as a key step to debug the problem in hand.

The motivation behind detecting the influence of third parties by end-hosts is two-fold. The third parties could also benefit from such tools. Some make changes to the stream of traffic and want to assess the effectiveness of their changes by testing whether their changes are noticeable by end-hosts or not. This is mainly because they hope their changes to the network flow to be undetectable by the end-hosts. For instance, they can claim that performance is not heavily affected by security augmentation to the stream, if the end-hosts are unable to sense the change. Conversely, a stealth MITM attacker or an unauthorized eavesdropper aims to remain undetected by end-hosts throughout the attack.

1.3 Definitions

In this section, we define technical terminologies that we will be referring to throughout this document.

Definition 1.3.0

An *end-host* is one of the two communicating parties in the network. An end-host is capable of observing everything happening on its own machine (e.g., how its congestion control mechanism behaves).

Definition 1.3.1

A *network flow (or traffic stream)* between two end-hosts is uniquely identified by the 5-tuple of the source and destination addresses, port numbers, and the transport protocol type [16].

Definition 1.3.2

A *middlebox* is defined as any intermediary device performing functions other than the normal, standard functions of an IP router on the datagram path between a source host and destination host [6].

Definition 1.3.3

A *translucent middlebox*[1] for a network flow between two end-hosts is an intermediary with the following characteristics:

1. The end-hosts receive exactly the same payload that the other party has sent. This means that the third party either does not modify the payload of any packets or undoes his changes before the packets get to the receiver[2].
2. The intention for influencing the network flow is more than just packet routing.
3. All packets in the network flow between the end-hosts go through the middlebox.

1.4 Middleboxes' Applications Taxonomy

In order to devise a non-technical taxonomy of this class of middleboxes, we observe the underlying principle behind why they exist along with their objectives rather than the technical characteristics of what they do to network streams. Figure 1.1 summarizes this taxonomy. In the remainder of this section we discuss each of the proposed categories.

1.4.1 Performance

Performance Enhancing Proxies (PEPs) [4] are broad range of examples of such third parties for network performance improvement. PEPs are used in satellite communications (TCP Split [11] and TCP Spoofing [8]), as well as in mobile network (ITCP [2], MTCP [23], and M-TCP [5]). PEPs are also used to cope with asymmetric links (ACK Filtering [3] and ACK Decimation [15]). These proxies are sometimes used to improve throughput of low bandwidth links by implementing various techniques such as compression or coding.

[1] From now on in this document we refer to translucent middleboxes as middleboxes.

[2] Note that middleboxes that drop packets are also included in this definition, whether the middlebox drops packets intentionally or unintentionally (e.g., as a result of saturated buffer queues).

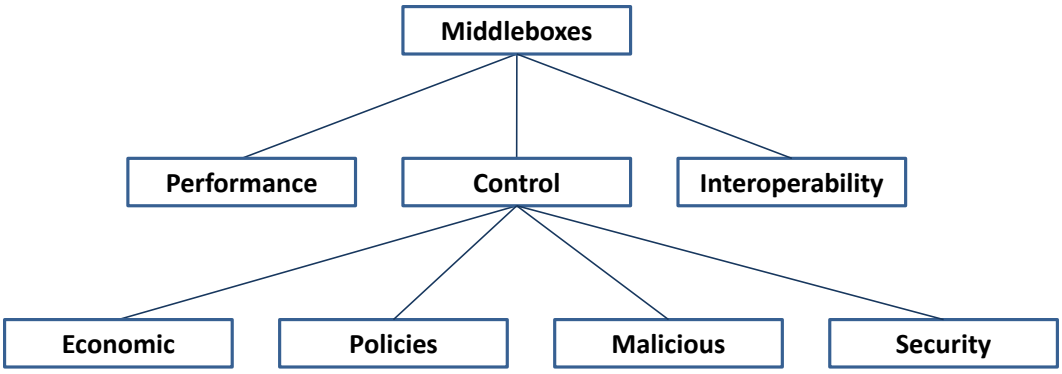

FIGURE 1.1: A non-technical taxonomy for middleboxes

Transparent web proxies [17] are another example of such third parties used to improve network performance.

1.4.2 Control

1.4.2.1 Economic

ISPs use a variety of tools (or a combination of such tools) for bandwidth management such as traffic shaping, policing, capping, throttling, classifying and conditioning. Another example is network address translators (NAT) which are widely used for the purpose of IP address management. Intermediaries also may impose proactive packet dropping to improve energy efficiency in wireless networks [7].

1.4.2.2 Policies

Government or companies may employ various techniques to control the incoming and outgoing traffic as well as the traffic flows within their network. Internet censorship, firewalls, and wire-tapping are well-known examples of deployment of such third parties. These organizations are also capable of controlling network applications using various techniques such as network dissuasion [13] and traffic classification. They can also specify regions of the Internet or countries they wish their network traffic to avoid [9].

1.4.2.3 Malicious

A man-in-the-middle (MITM) does not always modify the content of the packets. The attacker may simply eavesdrop on the traffic, and thus attack the confidentiality aspect of it. The attacker is also capable of causing service denial by performing either a high volume DoS attack or clever low volume DoS attacks (e.g., Shrew attacks [10]). Furthermore, Malicious selective packet dropping of critical network messages in wireless ad hoc networks can potentially paralyse the network by partitioning its topology [12].

In a different class of attacks, called delay attacks, a malicious attacker in the middle deliberately delays the transmission of time synchronization messages to magnify the offset between the time of a malicious node and the actual time [20].

1.4.2.4 Security

Practical examples of such influence by third parties are firewalls and VPNs (or more generally adding secrecy to the communication by encrypting the traffic flow). Packet marking (e.g., for IP traceback [21]) is another example of third party influences for security matters.

In our prior work [14], we introduced a third party intermediary, Personal Security Device, which is a portable device to improve security for mobile medical systems. The device we developed requires no changes to either the medical device or its monitoring software. The personal security device is designed to seem transparent to both parties so it could offer protection for millions of existing devices.

1.4.3 Interoperability

Sometimes third parties modify a network flow or a subset of packets in a particular flow for interoperability purposes. Tunnelling for various reasons such as IPv6 and Mbone (Multicast backbone), and IP fragmentation are examples of this kind.

References

[1] Enterprise network and data security spending shows remarkable resilience.

[2] A. Bakre and BR Badrinath. I-tcp: Indirect tcp for mobile hosts. In *Distributed Computing Systems, 1995., Proceedings of the 15th International Conference on*, pages 136–143. IEEE, 1995.

[3] H. Balakrishnan, V.N. Padmanabhan, and R.H. Katz. The effects of asymmetry on tcp performance. *Mobile Networks and Applications*, 4(3):219–241, 1999.

[4] J. BORDER, M. KOJO, J. GRINER, et al. Rfc3135. *Performance enhancing proxies intended to mitigate link-related degradations*, 2001.

[5] Kevin Brown and Suresh Singh. M-tcp: Tcp for mobile cellular networks. *SIGCOMM Comput. Commun. Rev.*, 27(5):19–43, October 1997.

[6] Brian Carpenter and Scott Brim. Middleboxes: Taxonomy and issues. Technical report, RFC 3234, February, 2002.

[7] W. Chen, U. Mitra, and M.J. Neely. Packet dropping algorithms for energy savings. In *Information Theory, 2006 IEEE International Symposium on*, pages 227–231. IEEE, 2006.

[8] J. Ishac and M. Allman. On the performance of TCP spoofing in satellite networks. In *Military Communications Conference, 2001. MILCOM 2001. Communications for Network-Centric Operations: Creating the Information Force. IEEE*, volume 1, pages 700–704. IEEE, 2001.

[9] E. Kline and P. Reiher. Securing data through avoidance routing. In *Proceedings of the 2009 workshop on New security paradigms workshop*, pages 115–124. ACM, 2009.

[10] Aleksandar Kuzmanovic and Edward W. Knightly. Low-rate tcp-targeted denial of service attacks: the shrew vs. the mice and elephants. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '03, pages 75–86, New York, NY, USA, 2003. ACM.

[11] M. Luglio, M.Y. Sanadidi, M. Gerla, and J. Stepanek. On-board satellite Split TCP proxy. *Selected Areas in Communications, IEEE Journal on*, 22(2):362–370, 2004.

[12] A.T. Mzrak, S. Savage, and K. Marzullo. Detecting malicious packet losses. *Parallel and Distributed Systems, IEEE Transactions on*, 20(2):191–206, 2009.

[13] Vahab Pournaghshband, Leonard Kleinrock, Peter L. Reiher, and Alexander Afanasyev. Controlling applications by managing network characteristics. In *IEEE International Conference on Communications (ICC)*, 2012.

[14] Vahab Pournaghshband, Majid Sarrafzadeh, and Peter L. Reiher. Securing legacy mobile medical devices. In *3rd International Conference on Wireless Mobile Communication and Healthcare*, MobiHealth '12, 2012.

[15] Segura R. Asymmetric networking techniques for hybrid satellite communications. *NATO Technical Note 810*, 2000.

[16] J. Rajahalme, S. Amante, S. Jiang, and B. Carpenter. Rfc6437: Ipv6 flow label specification. 2011.

[17] M. Shapiro. Structure and encapsulation in distributed computing systems: the proxy principle. In *The 6th International Conference on Distributed Computing Systems*, 1986.

[18] Justine Sherry, Shaddi Hasan, Colin Scott, Arvind Krishnamurthy, Sylvia Ratnasamy, and Vyas Sekar. Making middleboxes someone else's problem: network processing as a cloud service. *ACM SIGCOMM Computer Communication Review*, 42(4):13–24, 2012.

[19] Justine Sherry, Sylvia Ratnasamy, and Justine Sherry At. A survey of enterprise middlebox deployments. 2012.

[20] H. Song, S. Zhu, and G. Cao. Attack-resilient time synchronization for wireless sensor networks. *Ad Hoc Networks*, 5(1):112–125, 2007.

[21] S. Vincent and J.I.J. Raja. A survey of ip traceback mechanisms to overcome denial-of-service attacks. In *Proceedings of the 12th international conference on Networking, VLSI and signal processing*, pages 93–98. World Scientific and Engineering Academy and Society (WSEAS), 2010.

[22] Zhaoguang Wang, Zhiyun Qian, Qiang Xu, Zhuoqing Mao, and Ming Zhang. An untold story of middleboxes in cellular networks. *ACM SIGCOMM Computer Communication Review*, 41(4):374–385, 2011.

[23] R. Yavatkar and N. Bhagawat. Improving end-to-end performance of tcp over mobile internetworks. In *Mobile Computing Systems and Applications, 1994. WMCSA 1994. First Workshop on*, pages 146–152. IEEE, 1994.