Enterprise Environment Visiblity

Discussion Paper

September 3rd, 2020

CENTER FOR CYBERSECURITY
POLICY AND LAW

Introduction

Evolution in technology often creates challenges as governments and industry seek to balance the tradeoffs between security, privacy, visibility, and efficiency. When done correctly, these four elements combine to enable products and services that encourage innovation, keep users and organizations safe, and meet regulatory and legal requirements. However, when these are not in balance, potential issues can arise.

Once such technological evolution has been with encryption protocols such as Transport Layer Security ("TLS") 1.3. TLS has been a foundational element of network security for many years, helping to ensure that information passing between various types of end-points and applications is resistant to unauthorized tampering or exposure. This has provided a level of privacy and assurance for the public Internet that has helped to spur economic growth, information sharing, and countless other benefits to global society.

The same protocols designed to address security and privacy across the public Internet are also used in enterprise environments. As a result, the evolution of security technology can sometimes negatively impact an organization. For example, it can limit an organization's visibility into its own systems and networks. In the case of TLS 1.3, new features of the protocol are making traditional means of gaining visibility into network traffic much more difficult. This loss of visibility reduces the effectiveness of tools and methods which in turn limits their usefulness as security and troubleshooting capabilities. Ultimately, this means that enterprises are finding that previous investments in technology that have become cornerstones of their approach to network and system management are at risk of becoming ineffective.

This paper does not pretend to capture all aspects of this complex issue, nor does it offer definitive solutions. Rather, the purpose is to inform ongoing discussion.

Background

To help further the discussion around this important topic, on September 24th, 2019, The Center for Cybersecurity Policy and Law ("Center") hosted a multi-stakeholder meeting to discuss the implications of evolving protocols on the normal operations of both public and private sector enterprises. The goals of the workshop were to:

- arrive at consensus around a problem statement;
- identify potential near- and long-term solutions; and
- discuss next steps.

Workshop participants represented a cross-section of companies from multiple sectors, including financial services, healthcare, and telecommunications, as well as government agencies from the United States and the United Kingdom. Participants also included equipment and software vendors and members of civil society who provided insight into the concerns that have driven some of the protocol evolutions. Importantly, the participants also represented both operational leaders who have the day-to-day responsibility to manage their networks, as well as policy and management leadership. These diverse viewpoints were essential to revealing and understanding how the loss of visibility is impacting the business and mission.[1]

Enterprise Environments versus the Public Internet

First and foremost, it is important draw a distinction between enterprise environments, which represent the systems and networks operated by an organization in support of their business and mission, and the public Internet.

It is worth noting that this distinction does not mean that the two are mutually exclusive from either a technical or policy perspective. An enterprise that wants to maintain visibility within their networks is equally motivated to ensure that their information is protected from eavesdropping when it traverses the Internet, which it frequently does. Securing external connections with business partners, members of the supply-chain, and customers is essential to maintaining trust and strong encryption protocols enable that.

Similarly, customers who want their information protected on the public Internet also want the enterprises they entrust with their often highly sensitive information to be able to secure and maintain the infrastructure that supports the services on which they depend. Frequently, this includes data protected by law or regulation (e.g. personally identifiable information; personal health information), or proprietary information such as intellectual property, financial records, and so forth. As the custodians of this information, enterprises are under various obligations to protect what is in their care and maintain critical services to the maximum extent possible. This is necessary to ensure the trust of their employees, partners and customers and to avoid regulatory or legal action.

[1] A summary report can be found at https://centerforcybersecuritypolicy.org/enterprise-data-center-transparency-and-security-initiative

Legal Requirements and Industry Best Practices

Different regulators have established expectations for both data protection and enterprise security. A single entity or set of data may be required to comply with multiple regulations for different regulatory authorities within a single jurisdiction. While regulators increasingly expect entities to utilize consensus-based standards and controls to demonstrate compliance, bespoke and conflicting regulatory expectations still exist and impose burdens and technical challenges on effective data protection and enterprise security. Below are some examples, although this is far from comprehensive and meant to be illustrative.

The Federal Trade Commission ("FTC") has applied its statutory authority to enjoin "unfair or deceptive" business practices under the FTC Act[2] to enforce data security practices, and has brought enforcement actions under its unfairness authority against companies for engaging in practices that the FTC believes present an unreasonable risk to the security of the personal information of employees, customers, and consumers.[3] The FTC's enforcement actions have repeatedly identified an alleged failure to appropriately monitor access to and activity on a company's networks as a factor in claiming that the company has not implemented reasonable data security practices. These enforcement actions have focused on practices such as the failure to implement an intrusion detection program, to monitor system logs, to monitor outbound transmissions, to monitor the network, to monitor internet traffic leaving networks and engage in data loss prevention, and to scan for unauthorized programs and file sharing applications.[4]

FTC guidance relating to network monitoring identifies a need for companies to monitor outgoing transmissions for unexpectedly large amounts of data being sent outside the network and to investigate such data transfers to make sure they are properly authorized.[5] The FTC has reiterated its position on the importance of monitoring a company's network for unauthorized programs in guidance that it published addressing security concerns raised by peer-to-peer applications.[6]

A handful of states have statutes (or implementing regulations) that impose specific requirements related to system monitoring and detection. Under Massachusetts[7] and Oregon[8] law, companies are required to "regularly" test and monitor the effectiveness of key controls, systems, and procedures to ensure the program is operating in a manner reasonably calculated to prevent unauthorized access to or use of personal information.

---

[2] 15 U.S.C. § 45(a).

[3] Although FTC complaints, consent orders, and settlements are not law, the FTC has been the most prevalent regulator with respect to articulating what an unreasonable data security practice is, and FTC settlements are instructive as to what a regulator may deem unreasonable.

[4] *In re Lookout Servs., Inc.*, No. C-4326, ¶¶ 7 (F.T.C. 2011) (Complaint); *In re Dave & Buster's, Inc.*, No. C-4291, ¶ 7 (F.T.C. 2010) (Complaint); *In re EPN, Inc. d/b/a Checknet, Inc.*, No. C-4370, ¶ 6 (F.T.C. 2012) (Complaint); *In re Cbr Sys., Inc.*, No. C-4400, ¶ 9 (F.T.C. 2013) (Complaint); *In re The TJX Cos. Inc.*, No. C-4227, ¶ 8 (F.T.C. 2008) (Complaint); *In re Franklin's Budget Car Sales, Inc. d/b/a Franklin Toyota/Scion*, No. C-4371 ¶ 8 (F.T.C. 2012); *In Re D-Link Corporation and D-Link Systems, Inc.*, No. 3:17-CV-00039-JD, ¶ 15 (N. D. Cal. January 5, 2017) (Complaint).

[5] *Id.*

[6] FTC, "*Peer-To-Peer File Sharing: A Guide for Business*," at 5, *available at* https://www.ftc.gov/system/files/documents/plain-language/bus46-peer-peer-file-sharing-guide-business.pdf.

[7] 201 CODE MASS. REGS. § 17.03(h).

[8] OR. REV. STAT. § 646A.622(2)(d)(B)(iv).

International Organization for Standardization (ISO) 27001 has also been widely used to benchmark and evaluate whether controls are "appropriate technical and organizational measures". ISO 27001 requires vulnerability scanning and management processes, network protection, segmentation, and controls, continuous monitoring processes, malicious code protection mechanisms, audit requirements, activities, and verification, and event logging and review.[9]

Some sectors, such as financial and healthcare have additional regulations that also speak to these issues. Ultimately, it is clear that regulators are paying attention to and getting increasingly specific on how organizations protect the data in their care. Determining the right balance or protections continues to evolve.

Problem Statement

Even when information is not protected by law and regulation, there is still a motivation to keep it safe from unauthorized access or exposure. The web sites individuals visit, the people they associate with, their vacations plans, etc. are all types of information that many people choose to keep to themselves or within a circle of trust that they control. Without strong encryption, the risk of that information being exposed in unanticipated or unauthorized ways increases.

These complex motivations can become unwieldy and ultimately must be addressed using a balanced approach at both the technical and policy levels. To bring some focus to the challenge at hand, the following notional three-part problem statement is provided:

- **Maintaining the confidentiality, integrity, and availability of information and services in the enterprise environment is essential to mission success.**
- **Evolving protocols can negatively impact an organization's ability to maintain visibility in the enterprise environment, including threat detection and mitigation, troubleshooting, and service delivery.**
- **The impacts can impede an organizations ability to respond to regulatory, contractual, or other requirements.**

Within the context of this statement, we can examine some potential ways forward.

Approaches to Visibility

Below are the three most common ways enterprises gain network visibility, and most organizations use more than one as part of their network security and management program.

Deep Packet Inspection (DPI)

Broadly speaking, DPI is a monitoring and management technique that examines information traversing a network and reports on, and makes decisions about, how it will be handled.

In certain cases, in-line decryption can be used for DPI. However, there are also situations where passive decryption is required, such as retrospective decryption of packet traces during forensic investigations of east-west traffic where in-line decryption is not economical.

End-point Monitoring/Logging

[9] ISO 27001 §§ A.12.2, A.12.4, A.12.6, A.12.7, A.13.

End-point capabilities, including logging and management agents, are invaluable tools for enterprises to manage and secure their system.

However, the frequently unstructured nature of many log messages and the likelihood of logs being disabled makes their efficient and effective use as a sole method of visibility impossible as a practical matter. Also, storing log data frequently comes with an associated cost that can become exorbitant for even moderately sized organizations.

End-point monitoring agents have increasingly limited privilege as: (a) more activity is locked down in applications; (b) the agent itself may be susceptible to compromise; (c) some devices in an enterprise environment may not have any end-point monitoring solution at all; and (d) resource constrained devices may not support a rich-enough agent, as is often the case with Internet of Things devices.

Encrypted Traffic Analysis

Techniques are available that can identify certain classes of threats or risks based on analysis of encrypted traffic. However, this type of analysis is largely heuristic based, which may not detect unknown attack vectors. Nor does this approach address the troubleshooting and performance monitoring use cases.

Use Cases

None of the approaches discussed above are complete by themselves, but each provides an important element that enterprises leverage for use cases where visibility is essential. At a high level, these can be categorized into operations, security, and compliance.

Operations

For the purposes of this paper, the term “operations” is used broadly to capture the wide range of manual and automated activities necessary for enterprise technology to function in support of its business and mission goals. Within that context, there are two related yet distinct areas that depend on visibility into network traffic: troubleshooting and performance monitoring.

Troubleshooting

Enterprise systems and networks are inherently complex and highly optimized to meet the various demands of the customers they are supporting. In many cases, service level agreements require near perfect availability. The interdependence between routers, switches, firewalls, middleware, operating systems, and countless other devices are all places where errors can occur and degrade service. When that happens, the ability for automated and manual response processes to diagnose and correct the issue quickly is essential.

Lack of visibility can hinder this process by masking information on the network that could identify the root of service issues and inform corrective action.

Performance Monitoring

Similarly, performance monitoring is an integral part of maintaining and advancing the high levels of service enterprises require. While some applications provide their own diagnostic information, the ability to see traffic across the network provides an independent level of assurance and can help to correlate reporting from different business units, applications, etc. This can be important to reporting on and being in compliance with service level agreements.

Lack of visibility may result in important metrics being unavailable or incomplete, negatively impacting an organization's ability to demonstrate that performance targets are being hit or improved to meet business goals.

Security

While encryption is clearly a foundational element of securing data-in-transit, the protections afforded by encryption protocols in combating the risk of compromise do not prevent information from being exposed in a myriad of other more common ways. Most data breaches and exposures occur at the end point level and not from malicious actors intercepting it in transit over the wire. Existing protocols already make such an attack method difficult and time consuming, particularly when end points are generally easier targets.

The indiscriminate use of encryption, rather than applying it in a risk-based manner, can make it easier for an attacker to hide their action in plain sight by increasing the number of encrypted channels through which actions can be taken or data exfiltrated.

Specifically, the loss of visibility impacts three areas that all enterprises depend on: intrusion detection, compromise investigation, and forensic investigation.

Intrusion Detection

In many types of breaches, attackers move laterally across the enterprise, going from one system to the next for a variety of reasons. For example, once an attacker obtains credentials with sufficient privileges, often through phishing or other common techniques, they can exploit existing enterprise applications to directly access confidential data that may be stored on a variety of other systems.

Lack of visibility would impact the ability of network defenders to see intrusion related activity that can be detected on the network including:

- The introduction of malware from outside the enterprise via e-mail or web site downloads.
- The establishment of a command and control channel to an external site.
- Reconnaissance such as port scans.
- Attempts to gain elevated privileges via systematic attempts to login to hosts.
- Lateral movement of malware from one host to another.
- Data staging before a bulk exfiltration event.

Compromise Investigation

Security response teams regularly receive indications of compromise that they must investigate, including validating that a breach did occur (i.e., not a false positive), determining the source of the breach, determining the depth and breadth of the breach, and establishing a plan for eradicating malware, web shells, backdoors, etc. Because most system compromises are accomplished via network connections, security response teams often analyze network communications to and from the system early in their investigations.

Lack of visibility impedes their ability to respond by depriving them of a key source of information.

Forensic Investigation

Visibility isn't just about real time activity. Subsequent to an incident, it is often prudent to conduct a root cause analysis to help better understand how the attack was successful and what can be done to reduce the risk of a similar attack in the future. Additionally, a forensic investigation may be desirable to ensure that the full scope of the attack is understood. The investigation would clarify that the attackers didn't leave behind tools that would aid them in the future or that they didn't exfiltrate more information than initially believed. Regardless, attackers have become very good at removing or altering logs in order to cover their tracks or misdirect response efforts and investigations.

Compliance

The push for privacy and security compliance in law and regulation globally has been significant over the last several years, albeit with mixed results. Driven by numerous breaches and data exposures, governments have been actively seeking to better protect the sensitive information of their constituents. Part of this has been focused on the importance of encryption as a mechanism to protect that information. In fact, in many jurisdictions, if a breached entity can demonstrate that information was properly encrypted and the keys not exposed, notification of impacted individuals may not be required.

Compliance is not limited to government regulation, and most organizations require compliance and reporting on security and performance controls as part of their business arrangements. Industry sectors may also have compliance regimes, such as Payment Card Industry's ("PCI") Digital Security Standard ("DSS")[10] within the financial sector.

Alternative Approaches to Enterprise Environment Visibility

Finding a balance between enabling enterprise environment visibility and securing the public Internet is possible but will likely require the development and/or approval of alternative methods that either extend existing methods or introduce new innovations. Regardless of what alternatives are eventually put into place, they should meet this set of notional criteria:

- Must be scalable.
- Must be relatively easy to implement/deploy.
- Must be protocol agnostic.
- Must be usable in real time and post packet capture.
- Must be effective for security, operations, and compliance purposes.
- Must be widely available and supported in mainstream commercial products and services.

This list of potential approaches to maintaining visibility in the enterprise reflects different levels of abstraction; some address visibility broadly while others are specific to the deployment of TLS 1.3. This should not be considered a definitive list, and further evaluation of the approaches is clearly warranted.

Overlays/microservice mesh

[10] https://www.pcisecuritystandards.org/

Service Mesh architectures such as ISTIO[11] help to control how different applications and services can communicate in container and other environments. For example, in Kubernetes, traffic to and from a specific pod is routed via a module called a sidecar, which provides a potential location to perform TLS offload and/or to deploy a man-in-the-middle proxy.

Key Retention

Key retention systems allow the key material generated in forward secrecy schemes such as Ephemeral Diffie-Hellman ("DHE") to be stored. The key retention system can retain the session key material for a short time, e.g., to support real-time decryption, or for a longer period to allow post-capture decryption later. True forward secrecy is achieved for a given time interval once the corresponding key material is permanently deleted.

Managed Diffie-Hellman

Rotation of static keys does not achieve per-session forward secrecy in the way that DHE does, but it can still provide forward secrecy after a given time interval. By increasing the frequency of key rotation, forward secrecy can be achieved after shorter time intervals. As with key retention, true forward secrecy is only achieved once the private keys used by both parties and the session keys are permanently deleted.

Enterprise Transport Security ("ETS") Standard

ETS is described in European Telecommunications Standards Institute ("ETSI") Technical Specification 103 523-3. It describes a profile for TLS 1.3 that uses static Diffie-Hellman keys. ETS includes requirements to notify the user of the TLS client, through either technical or procedural methods, that a static key is in use. ETS could be used as the basis for a Managed Diffie-Hellman framework.

RHRD

RHRD refers to a proposal made to the Internet Engineering Task Force ("IETF") during the TLS 1.3 discussions. It proposed a handshake extension that allows the client to opt-in to passive decryption of the session. The proposal can be found here: https://tools.ietf.org/html/draft-rhrd-tls-tls13-visibility-00

Multi-Party Computation ("MPC")

In classic Diffie-Hellman, only two parties are involved in the key agreement process. MPC allows more parties to be involved and therefore provides a means of sharing key material with an authorized decryption tool. One option would be to use an MPC framework with the pre-shared key ("PSK") capability in many encryption protocols, including TLS. At this time, this approach is largely in the conceptual phase and may not be a reasonable near-term solution.

Hardware Security Module ("HSM")

An HSM can perform the encryption handshake calculations on behalf of a server such that the static keys involved never need to leave the HSM, thus preventing them from being stolen. HSMs alone do not offer a solution to visibility challenges, but they can enhance the security of implementations of many solutions. In particular, there is discussion as to whether HSMs could be used to improve the security of a key retention or key rotation framework.

[11] https://istio.io/

Recommendations

There is considerably more to be done on this issue. While TLS 1.3 has been the primary subject of this paper, there is every reason to believe that as efforts to further secure the Internet continue, new protocols will be introduced, or existing ones updated, that will require consideration on how they impact enterprise environment visibility.

To keep the discussion moving, the Center recommends the following:

Research

Research has been repeatedly identified as a key area where progress could be made, both in terms of the viability of proposed technical solutions, as well as providing additional data points concerning the impact that evolving protocols can have and are having in enterprise environments. Some potential research efforts could include:

- Using existing government and private sector data sets to conduct various types of analysis to further refine understanding. Some questions that could be explored might include:
    - Do the current proposed approaches meet the principles discussed above?
    - How many actual, recent enterprise breaches have resulted from malicious actors decrypting data captured over the wire as opposed to compromising end points or stealing credentials?
    - How often, and in what ways, are malicious actors using encrypted channels once they have gained unauthorized access?
    - How often do network managers use TLS inspection to troubleshoot applications and detect and mitigate attacks?
- Define and establish a project within the National Cybersecurity Center of Excellence ("NCCoE")[12] at NIST to demonstrate that one or more potential visibility approaches is deployable and scalable in the enterprise environment within the context of one or more use cases.

Awareness and Education

There is concern that many organizations remain unaware of the challenges they will face as TLS 1.3 and similar protocols are promulgated. To that end, education and outreach efforts may be worthwhile.

Some ideas include:

- Draft white papers or other publications to increase CIO- and CISO-level awareness of the trend toward reduced visibility and its implications.
- Additional workshops to further define the challenge with different stakeholders and/or to further explore potential approaches.
- Increase awareness within protocol standards bodies in general by ensuring more stakeholder voices are present and heard during standards development.

[12] https://www.nccoe.nist.gov/

- Increase participation by technical staff (including protocol designers, engineers, and implementers) in the IETF and other standards bodies to champion enterprise use cases and help shape the evolution of the protocols.

Deployment Guidance

Additional deployment guidance would likely be beneficial for enterprises deploying TLS 1.3 within their environments.

- Develop guidance to help organizations manage TLS-based services as a critical asset and maintain cryptographic agility when migrating to TLS 1.3 or the Perfect Forward Secrecy with TLS 1.2 by enabling DHE key agreement.
- Develop tools and techniques to audit services for the unexpected deployment of DHE.
- Proactively determine where loss of visibility may impact the ability to troubleshoot problems.
- Ensure that TLS guidance does not unnecessarily become regulation.

Conclusion

Broadly speaking, encryption has positively impacted the growth of the internet by enabling organizations and individuals to protect information from unauthorized access. From digital commerce, to personal information, to national secrets, the use of encryption has been, and will continue to be, foundational and ubiquitous.

But as with any technology, encryption can be deployed in ways that cause more harm than good. Encryption alone cannot solve the myriad security challenges we are faced with. As such, it must be deployed in ways that are tied to risk and that don't prevent other important security measures from functioning as needed. Further, as policy makers struggle to strike the right balances within regulation, they need to understand the intricacies involved with encryption and not simply latch onto broad concepts as either "good" or "bad"; context matters.

Alternatives to current approaches exist and are being further developed. Backed by more research, pilots, and data gathering, it will be possible to achieve the benefits provided by intelligently applied encryption without curtailing the tools network defenders need.