

WHITE PAPER

What Do You Mean TLS 1.3 Might Degrad My Security?

How and why the new TLS 1.3 standard may affect security for better and for worse

The Disruption-Defense Conundrum

Transport Layer Security (TLS), formerly known as SSL, has become the de facto way of encrypting data in motion on networks. Unfortunately, several serious attacks have affected TLS over the past few years, and malware increasingly uses SSL/TLS sessions to hide, confident that security tools will neither inspect nor block its traffic. The very technology that makes the internet secure can become a significant threat vector. As the volume of encrypted traffic continues to grow, organizations become even more vulnerable to encrypted attacks, hidden command and control channels, and unauthorized data exfiltration exploits that go undetected. For this reason, the Internet Engineering Task Force (IETF) has voted to approve an updated version — TLS 1.3 — of the standard. Some cryptographers believe the new standard will be faster and more secure. Enterprises, on the other hand, are right to be concerned about the implementation and availability issues TLS 1.3 might cause. That is because TLS 1.3 has removed certain visibility that was widely deployed for threat identification in TLS 1.2.

Once again, InfoSec teams find themselves at the fulcrum of a delicate balancing act. On the one hand, encryption is moving toward ubiquity, but on the other hand, InfoSec teams need to be able to detect when threat actors use it too. What can you do? This whitepaper will delve into TLS, look at the security implications of TLS 1.3 and what you can do to prepare.

What is TLS?

TLS is the modern name for SSL (Secure Sockets Layer), although both terms are still used interchangeably, although calling it SSL is technically incorrect. TLS is a standard to secure communications between a client and server, but more generally between clients and applications that typically sit over a reliable transport layer, such as TCP, although there have been adaptations to UDP as well.

TLS provides three key services:

- **Confidentiality:** Ensuring that anyone intercepting the communications between the client and server cannot decipher that content.
- **Authentication:** Ensuring that a client is in fact talking to the server that the client thinks it is talking to. Optionally, the server can authenticate the client, but this is rare.
- **Integrity:** Ensuring that the messages and communication have not been corrupted or tampered with.

The confidentiality is ensured by leveraging symmetric cryptography, the keys of which are negotiated during a TLS handshake. The authenticity is established by using certificates, once again exchanged during the initial handshake, and maintained through the session using either HMAC (Hashed Message Authentication Codes) or AEAD (Authenticated Encryption with Associated Data), depending on the negotiated cipher suite. The integrity is ensured by using a Message Authentication Code (MAC).

Changes in TLS 1.3 and Their Implications

TLS 1.2 provided cipher suites that offered a choice of KEA, which meant you could use a non-PFS (perfect forward secrecy) cipher suite, typically RSA, to support passive interception. TLS 1.3 has removed static RSA and Diffie-Hellman cipher suites and only supports KEAs which use PFS.

TLS 1.3 has several changes that improve performance and security, while also eliminating several complexities and simplifying the protocol stack. However, there are implications for enterprises that use network security-based solutions for compliance, risk management, as well as threat hunting.

Table 1. TLS 1.3: A Summary

The Good	The Bad	The Ugly
Significantly reduced latency	The threat model that makes perfect forward secrecy highly desirable on the open internet is not relevant in closed data center environments.	Encrypting the session after the server hello makes certificates invisible, thus depriving the ability to perform valuable security analytics on the certificate.
Lots of simplification	MitM (man-in-the-middle) or agent-based approaches are undesirable in many environments, for architectural and performance reasons. Agentbased approaches are unavailable on some operating systems, which do not support userinstalled code.	The base requirements of enterprise threat detection seem to have been ignored or rejected by IETF.
Elimination of legacy features that are now considered undesirable	Workaround: Use TLS 1.2 in the data center and 1.3 outside.	

While there are many changes in TLS 1.3, here are two of the most important:

1. Reduced Latency

One of the key performance improvements in TLS 1.3 is speeding up the time it takes to negotiate protocol versions, cipher suites and authenticate the server. This is done during the initial handshake between the client and server. With TLS 1.2, this would typically take two round-trip times (2-RTT), as well as other protocol overhead. Simply put, a round trip time (RTT) is the time it takes for a client to send a message to the server and for the server to respond back to the client.

In TLS 1.3, this initial handshake has been cut down to 1-RTT. Furthermore, there is the ability for clients to “remember” previous sessions and resume the session without utilizing even 1-RTT. This would allow session resumption with a 0-RTT, making for far lower overhead, leading to faster time to connect to servers and load web pages, do transactions on the internet and in general be a more responsive browsing and internet experience.

Indeed, session setup latency is one of the big drivers for TLS 1.3 because it directly affects the end-user experience. However, there are some security implications associated with “resuming” a session with 0-RTT — potentially creating a window of opportunity for replay attacks. Because of this, many cryptographers aren’t convinced it’s a good idea, and some organizations have noted to Gigamon that they’re planning to disable this feature when they migrate their servers to TLS 1.3.

2. Only Forward Secrecy

Another major difference in TLS 1.3 is that the use of static RSA and Diffie-Hellman key exchange has been replaced with Ephemeral Diffie-Hellman, thereby providing forward secrecy. Forward secrecy means that if someone at some point in the future gains access to a server’s private key, they should not be able to decrypt all the past conversations if they have logged the session and stored it.

While RSA and Ephemeral Diffie-Hellman both generate unique keys per session, they do so in different ways. For RSA the client generates keying material, encrypts it with the server’s key from the certificate, and passes the encrypted keying material to the server. The server can decrypt the keying material using its private key, but imagine there is an entity that recorded data and stored it, maybe for years, and later obtained the secret key from the server. That entity could decrypt those key exchanges, and decrypt all of the historical data. This is because RSA lacks the forward secrecy property.

In addition to the use of Ephemeral Diffie-Hellman, all handshake exchanges between the client and server after the initial server hello are encrypted. This includes all the certificate data used in the handshake.

The Implications Matter

While that may sound great, the implications are significant. Many organizations today leverage network security appliances that need to inspect the data that is traversing their network.

This is necessary to:

- Identify malware and attempts to exfiltrate data
- Investigate malpractice, misbehavior, potential beaconing and other malicious communications

Active and Passive Mode Decryption

In some instances, content inspection is also used for troubleshooting, as well as for monitoring application performance. Devices that do this type of decryption may sit in a passive or out-of-band mode, or may be in-line to the network traffic either as a bump-in-the-wire device or as an active network element such as a firewall or Intrusion Prevention System (IPS).

Inline TLS decryption devices typically function as a man-in-the-middle, intercept a client’s connection to an application server and set up a separate connection to the client and to the server, each of which terminate on the man-in-the-middle device.

Passive mode deployments are common in many large enterprises, particularly those that have strong regulatory and compliance requirements. When deployed in a passive mode, traffic decryption is possible when:

- Using the RSA key exchange
- The application server’s private keys are shared with the decryption device
- Traffic is typically in-bound to the application server

With TLS 1.3, this passive mode decryption will no longer be possible since the RSA key exchange has been removed. This means that organizations that were leveraging passive mode devices that decrypted content, based on policies, will no longer be able to do this for threat hunting or regulatory compliance.

This is a huge change that will force many industries to rearchitect how they do security, compliance or monitoring. It could also slow for the adoption of TLS 1.3. Note that using inline man-in-the-middle device solutions for decryption are still possible — for now — and organizations may be able to take advantage of these types of solutions to enhance their compliance and security posture. In some situations, MiTM TLS inspection may introduce challenges, as it requires architectural changes, adds latency and cost, and may bottleneck performance. There is also the practical limitation that you can deploy man-in-the-middle devices only in so many places, since almost all such solutions do introduce some latency.

Hence, maintaining coverage, for example, for internal/lateral (East-West) traffic, would be difficult to achieve. Inspecting and monitoring East-West traffic is important and can indicate good behavior or a malicious attack. Because East-West traffic travels inside an organization’s network, it does not touch the perimeter where controls are typically deployed. With TLS 1.3, the potential risk is that as East-West traffic becomes more and more difficult to decrypt, organizations will stop looking at it — and that could prove risky. While there were proposals into the IETF that provided opt-in passive monitoring capabilities for TLS 1.3, specifically targeting data center environments, the IETF chose not to proceed with those. It seemed like the needs of these customers was disregarded.

Another area of potential impact is that with TLS 1.3, all packets in the handshake after the initial server hello are encrypted. This includes the server certificates. Network security solutions that relied on understanding the information in the TLS handshake — for example, examining the server certificate to help identify self-signed certificates or other anomalous situations — will no longer be able to do so.

Server-side authentication, which provided invaluable security analytics data without decryption under TLS 1.2 and earlier, has disappeared. Once again, this impacts passive-mode network security solutions. Inline solutions may still be able to decrypt the traffic and examine certificate data or other data in the TLS handshake, for anomalies or threat hunting. Again, the limitations outlined above, i.e., the number of places where this can be done, will be limited due to practical considerations, but at least for internet traffic, this may be a possibility.

TLS 1.3 Adoption

From an adoption and timing perspective, while some browsers today are already shipping with TLS 1.3 capability, the internet-wide and industry-wide rollout of TLS 1.3 will take time. There are still implementations of TLS 1.0 and 1.1 out there today, and indeed SSLv3 and even SSLv2 are occasionally seen in enterprise environments. While some browsers have supported TLS 1.3 quickly, organizations impacted by the above changes may be slower to migrate their application servers to TLS 1.3, even when it becomes available. They first need to understand how to re-adjust their compliance and security stack. This also disregards IOT and OT devices, which may never upgrade to TLS 1.3, for a variety of reasons.

The shift to TLS 1.3 may also be hampered by the presence of man-in-the-middle network devices that perform TLS decryption. These devices implement the TLS stack themselves and serve as a TLS proxy to intercept a client's TLS connection to the server. Until these devices also migrate, the adoption of 1.3 may be limited. Note that TLS 1.3 will allow a connection to down-negotiate to TLS 1.2 when either side does not support TLS 1.3, and this applies to man-in-the-middle devices as well.

Preparing for TLS 1.3

While the new standard has been in the works for some years, the final form was not released until August 2018. For those organizations on the fence about what they want to decrypt, this could be a forcing function as there is so much complexity and so many different places where data resides. This is a new opportunity to begin looking for threats under the lens of, "If I have to pick certain parts of communications to decrypt, what would they be?" Much of this will be driven by regulations, brand reputation or data volume.

You can always do business; the question is "at what cost?" Can you put something inline that will terminate the connections and not add latency or risk to your operations? Or, do you have to put termination inline in hundreds of places and hire new people to make sure it all works? Or, could technology come along that enables you to do this in a way that still maintains some level of privacy and security while also allowing you to meet required controls?

For instance, regulated industries like healthcare and financial services, which have to comply with HIPAA, PCI-DSS or GDPR, may face certain challenges when moving to TLS 1.3 if they have controls that say, "None of this data will have X, Y, or Z in it" or "This data will never leave this confine and we can prove it by inspecting it." In order to prove compliance with

those controls, they must look inside the encrypted traffic. However, if their infrastructure can't see traffic or is not set up to be inline with everything that is out of band in their PCI-DSS, they can't show that their controls are working. If they're out of compliance, therefore, they might also be out of business.

For others, perhaps newer companies, this might not be a huge concern. However, for institutions contending with hefty investments in legacy infrastructure, the change could be difficult to manage. It also brings up the question: Is it reasonable to introduce new costs and potential reliability risks in the name of better security? Security usually takes a back seat to operational efficiency until a compromise occurs. Then the lament is, "We could have avoided this."

In summary, as organizations continue to move their applications and workflows to the cloud, the volume of encrypted traffic is increasing and companies are upgrading to faster networks to keep up. TLS 1.3 has significant changes, improvements and some simplifications that can improve the security and performance of internet communications and transactions. However, the rollout of 1.3 may take some time and organizations should carefully understand their compliance and regulatory requirements, and the role of their network security solutions while taking a phased migration approach.

How Gigamon Can Help

The Gigamon Security Delivery Platform's SSL/TLS Decryption solution, with inline capabilities, brings visibility into encrypted data and is the first solution to run on high-speed networks, enabling security operations to take a unique architectural approach of a decryption zone to solve the problem of SSL/TLS decryption. In a decryption zone, SSL/TLS traffic is decrypted once and fed to multiple security and operational tools for analysis, thereby eliminating unnecessary and repetitive cycles of decryption and reencryption within the infrastructure — assuring security, eliminating unnecessary spend and increasing ROI.

With this approach, customers of the Gigamon Security Delivery Platform saw more than a 50 percent reduction in security costs and 153 percent return on investment, according to Forrester Total Economic Impact™ Study commissioned by Gigamon in 2016.

Gigamon is committed to supporting TLS 1.3 in its SSL Decryption solution, as the new standard is adopted within internet and enterprise networks.

How Gigamon Helps SecOps Teams



Figure 1: How Gigamon Helps SecOps Teams

Next Steps

Learn more about how to strengthen your security with Gigamon by visiting [Gigamon.com](https://www.gigamon.com).

About Gigamon

Gigamon is the first company to deliver unified network visibility and analytics on all data-in-transit, from raw packets to apps, across physical, virtual and cloud infrastructure. We aggregate, transform and analyze network traffic to solve for critical performance and security needs, including rapid threat detection and response, freeing your organization to drive digital innovation. In short, we enable you to run fast, stay secure and innovate. Gigamon has been awarded over 75 technology patents and enjoys industry-leading customer satisfaction with more than 3,000 organizations, including 80 percent of the Fortune 100. Headquartered in Silicon Valley, Gigamon operates globally. For the full story on how Gigamon can help you, please visit www.gigamon.com.