THALES E-SECURITY

Data Transmission Security:

Securing the Blurry Line between Classified and Unclassified Data

by Joseph Warren, Network Security Product Line Manager

Introduction

Now more than ever, data transmissions are an increasingly prevalent part of every military mission. From Headquarters to Tactical Operations Centers to the individual warfighter, the amount of data is massive and clear distinctions between classified and unclassified data transmission methods are progressively blurred. Top Secret and Secret classifications of data are well established and the methods for transporting these data types are known. But new methods are required to protect classified data as a result of the sheer volume of data, the time sensitive nature of the data, the number of data recipients, and the constraints associated with classified encryption devices. Debate over commercial encryption and double-commercial encryption methods for these new data classes continues. This paper will present a data transmission security strategy that can protect sensitive and unclassified data while preserving the integrity of established classified standards with minimal impact on network overhead and performance.

What Are The Threats?

The quality of a data transmission security strategy must be measured in terms of a total solution. There are many components required to successfully send a packet over a public and/or private network infrastructure and all of these components must be protected against threats. And just to be clear, private infrastructures are not inherently any more secure than a public infrastructure. For example, a Virtual Private Network (VPN) is merely a virtual segmentation of a larger network. The word “private” is misleading in the sense that privacy can only be obtained if unwanted guests cannot see or get inside. There is a massive difference between the terms “private network” and “secure network”. So what exactly are the vulnerabilities that should be considered when looking to secure the transmission of data?

Direct Threats

Data transmission threats come in many different forms. The threats are obvious when there is no security at all. When data in motion is unsecure, an eavesdropper could, at a minimum, intercept, read, manipulate, and impersonate the data stream.

Indirect Threats

Indirect threats are almost always available to an eavesdropper, even when the data is thought to be secure. For example, encryption helps to secure data by making at least part of the stream unreadable to the eavesdropper. But if the source and destination addresses are in the clear, end users might be identifiable. If the transmission is not properly secured, packet replay could occur, resulting in potential database corruption and other forms of malice. Even the mere structure and size of encrypted packets can give away useful clues to a sophisticated listener. These are all examples of indirect threats that must be accounted for when securing data in motion.

Encryption in itself does not constitute data transmission security

To begin with, encryption can secure data but does not necessarily secure the transmission of the data. Data security and transmission security should be thought of almost as two exclusively separate topics that work together to solve the holistic problem of secure data transmission. Since encryption provides the security of the data, a best practices encryption approach should be applied. Data encryption is in itself, a topic for a scrutinized discussion. Take for example, the lifecycle and management of the keys used to encrypt the data. A device might implement an internationally recognized algorithm such as AES256, but if the key material is weak, the encryption is easily debunked.

Key Entropy

The security of a key begins with the randomness of the numbers generated to create keys (key entropy). Imagine if you encrypted a message by simply changing each letter of the alphabet to the next letter, substituting an “A” for a “B” and a “B” for a “C” and so on. Let us call this “ABC encryption” as an example of a simplistic implementation. With ABC encryption, the key would not be very difficult to break because it is not very random and certainly not long and complicated. While this is a simplistic example, it should be considered that with powerful

computing capabilities that exist today, vulnerabilities in even extremely complex key material could easily be exploited. Use of hardware based Random Number Generators as a source to produce completely random numbers is well recognized as being superior to pseudo random number generators used in software-based implementations.

Key Storage

Now imagine the over exaggerated example of writing down the key onto a piece of paper and leaving it on top of a desk. This would completely undermine the security of everything from our simplistic ABC encryption example all the way to highly secure government encryption methodologies. In order to protect keys from prying eyes, key material must be stored within a hardware boundary that, if tampered with, immediately zeroes out the data and completely renders the device from operating at all. Secret key material should never be seen by humans nor should it be evaluated by intrusive mechanical and/or electromechanical methods.

Key Lifecycle

In addition to randomness and storage of key material, keys should be changed often. In our ABC encryption model, the encryption could be broken quite quickly, facilitated by the use of one key (each letter is shifted 1 character). But what if the key were changed after each word? Word one would have letters shifted by one character, word two might have letters shifted by two characters, word three might have letters shifted by three characters, and so on. Granted, this would still not be a very sophisticated method of encryption however, it does exemplify how much greater the security becomes simply by changing the key material often.

Key Transport

So far this paper has established that key entropy, key storage, and key life are all major components of a solid encryption solution. It should also be recognized that the transport of key material between trusted devices is yet another essential requirement. If deficiencies within the transport methods exist, serious consequences can result. While widely accepted

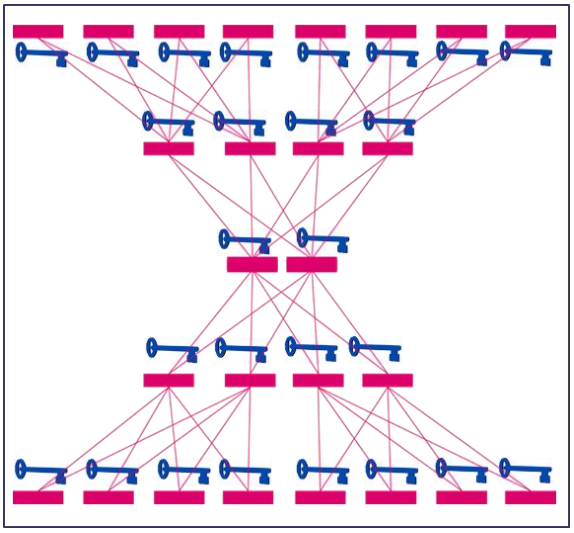

methods such as Diffie Hellman key exchanges ensure key material secrets are kept safe, consideration should be made for efficient transport of the key material. For example, in a large network where keys are changed quite frequently, a situation could arise where more time is required to key devices than is allotted for data transport. The result would be a secure key management system that cannot pass data. Using a group key management scheme that leverages multicast protocols to deliver key material through a secure channel would help to ensure that even large-scale networks can be keyed efficiently and without interruption of service. Careful consideration for the methods of key and data transport are essential to both the security and performance of the data in motion security puzzle.

Susceptibilities of Encrypted Data in Motion

Encrypted data can be considered quite secure when it is stationary. However, the transport of encrypted data through a network (public and/or private) brings a whole new set of vulnerabilities to the security puzzle. If we examine the OSI Model (the 7 layers of networking), it is clear that as you move up in layers, more capabilities are provided. Whenever more capabilities are added, more vulnerabilities become possible. Data, even encrypted data, has vulnerabilities when it is set into motion.

Man in the Middle Attacks

Even if a solid encryption and key lifecycle management solution has been implemented, once the encrypted data gets set into motion, prying eyes could view and intercept the encrypted data. If the encrypted data does not have appropriate transport security mechanisms in place (such as Galois Counter Mode and Frame Check Sequence), data integrity can be compromised and packet replay attacks can undermine security. Both of these are examples of attacks that can be accomplished at the transport layer without the interceptor ever being able to decipher the encrypted data. A sophisticated data in motion encryption solution that does not provide for data integrity can easily facilitate database corruption and other man in the middle attacks. Be wary of data in motion encryption solutions that claim "low to no" overhead as a method of transport efficiency. A few less bytes of overhead usually translates into less transport security.

Data Analysis

Additional information can be ascertained by eavesdroppers even if data is properly encrypted and appropriate transport security mechanisms are in place. These vulnerabilities exist in the examination of encrypted packets as they pass through the infrastructure. Eavesdroppers might not be able to decipher the data and they might not be able to change or replay the data, but they can glean plenty of information from analysis. An eavesdropper might notice that encrypted jumbo frames move out of a site at midnight each evening or that small packets begin from that same location each day at 0900. That eavesdropper might determine that the small packets represent encrypted voice streams (IP Phones) or that the jumbo frames represent a data center backup each evening. The eavesdropper might even be able to tell the destination address of the packets indicating the location of the data center backup. All of this information could be learned without ever knowing the content of the encrypted data. Providing a method for obscuring the data through implementation of Traffic Flow Security (discussed in the *Throughput* section) is yet another essential component of a total data in motion security solution. Traffic Flow Security should be implemented to mask data patterns and to prevent and eliminate data analysis.

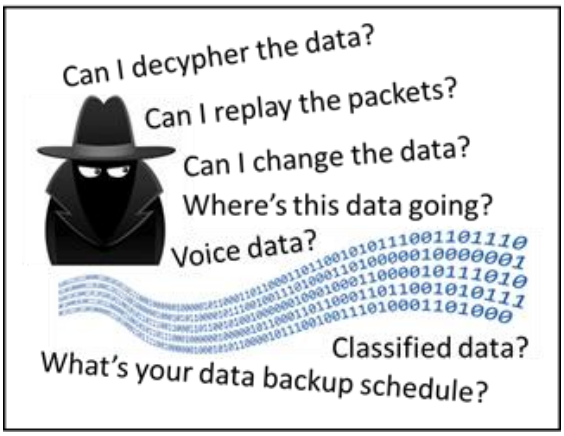

Efficiency of Encrypted Data in Motion

Network efficiency is usually measured in terms of the speed in which it takes to deliver data across a network. It is a confusing measurement because of the many different types of networks, the amount of traffic on the network, the distance and number of routers/switches between sites, the types of transport security mechanisms in place, etc. But for purposes of this focused data in motion security discussion, efficiency should be measured in terms of speed of packet delivery between encryption devices. This speed should be calculated with both encryption and total transport security mechanisms in place.

Latency

Latency, for the purposes of this discussion, is the time that it takes to encrypt a packet and pass it on to a paired device where it is then decrypted. If encryption is performed in hardware, such as a Field Programmable Gate Array (FPGA), the encryption and decryption process is done at near instantaneous speeds. Most encryption manufacturers leveraging FPGAs boast latency figures of around 4µsec at speeds of 10Gbps. Since the encryption calculations are done at the hardware level, the 4µsec latency figure is consistent at different packet sizes. This is important because different data types have different packet sizes. Voice packets, for example, are usually small in size while jumbo frames can be more than 1000 times larger. So why is packet size relevant to the discussion of

latency? It is important because network data does not consist of only voice, or only video, or only data, or only jumbo frames. It is an ever-changing mix of information where the network requirements of tomorrow could be quite different from the network requirements of today. Encryption solutions that leverage microprocessors to encrypt and decrypt data (such as those found in many routers and switches), have different latency figures based on packet size. Millions of small packets at 10Gbps place a huge demand on microprocessors. Likewise, large jumbo frames take a long time to calculate the hash. The importance of performance consistency derived from an FPGA implementation cannot be stressed enough. Consistency of encrypted performance regardless of packet size is extremely relevant to the discussion of efficiency.

Throughput

Throughput efficiency is measured in terms of how much data can be passed to another device divided by the time it takes to complete the transfer (Throughput = Data/Time). Throughput results (in a controlled test environment) are greatly influenced by the size of the data packet, the transport layer, and the latency of the encryption devices. For a moment, let us eliminate the latency variable since that is unique to each encryption device. Packet size and transport layer affects throughput because of the amount of overhead associated with each packet. Source and destination addresses, initialization vectors, counters, frame checks, frame gaps, and even MPLS tags are all examples of overhead that could be found within each frame and all are in addition to the data itself. The overhead is constant regardless if you are sending small amounts of data, or very large jumbo frames. So it is easy to see that the smaller the frame size and the more transport overhead you have, the greater the ratio of overhead to data is. In fact, it is entirely possible to become inefficient to the extent that the amount of overhead is greater than the amount of data being transported on a per frame basis.

With efficiency being the goal, it is possible to create constant frame sizes that are limited only by the frame sizes that the network can handle. We will call these manipulated frames

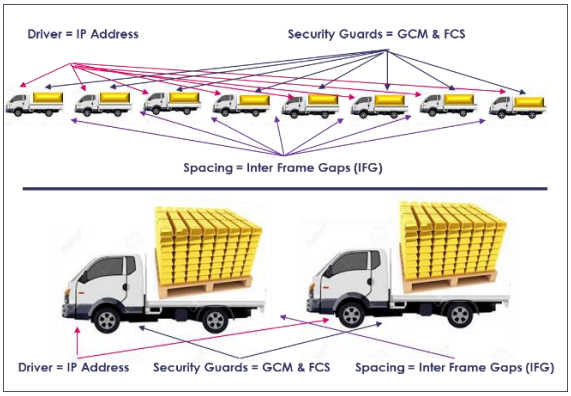

Figure 1 – Overhead and its effect on efficiency

"containers". When these containers are maximized, it becomes possible to insert multiple frames into the container. This eliminates the need for individual frame overhead constraints (see Figure 1). For example, by creating containers of around 3000 bytes, one could take about forty-eight 64 byte packets and place them into a single frame. When properly implemented, the GCM, FCS and IFG overhead can be completely eliminated for individual packets and implemented only once in the temporary 3000 byte container. The result would be a more efficient way of transporting data while maintaining the security of the transport layer. If we use the analogy of moving millions of bars of gold from one place to another, there are many ways to accomplish this. We could move each bar of gold one bar at a time in a small truck. Each truck would require a driver (IP Address), each truck would require security guards (GCM and FCS), and each truck would require space between each other (IFG) so they would not crash into one another. But if the gold were moved at 50, 100 or 1000 bars at a time, think of how much more efficiently the gold would get to its final destination. Fewer trucks means fewer guards, fewer drivers, and fewer spaces in between. In fact, test data shows that using these types of containers produces up to 98% network efficiency for encrypted data, including all requirements for transport security as described in this paper. The end result is a completely secured stream of encrypted data with only a marginal 2% impact on network performance. All of the benefits of encrypted traffic with complete transport security can be achieved at nearly no cost to network performance, regardless if voice, video, data, or

jumbo frames are sent. In the event that no data is being sent, random data can be generated to completely obscure the line from any traffic patterns whatsoever. High levels of security (including prevention of traffic analysis), consistency in performance, and nearly immeasurable impact to throughput can finally be achieved.

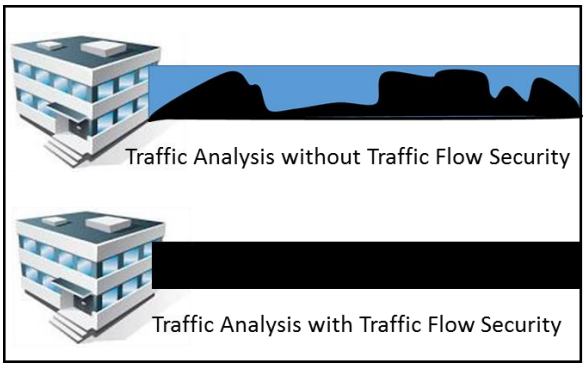

Transport Security and Complex Data Classification Policy Issues

Now that it is established that full data in motion security for unclassified data can be accomplished at 98% efficiency, a baseline can be set for solving data policy issues. Data is now a critical and embedded part of every warfighter's mission. The ability to parse classifications should not be constrained by the inability to protect data in a logical manner. Top Secret and Secret classifications of data are known and well established. The mechanisms for transporting these data classifications are also well established. But new classifications of data continue to emerge and debate over commercial encryption and double-commercial encryption for certain data classes continues.

The Data Classification Debate

The debate over data classifications and the methods for transporting that data are decisions that are best left to the governments and organizations tasked with protecting the data types. But generally speaking, the debate exists because of the sheer quantity of data being passed. And data is sometimes only relevant for a given period of time. For example, it might be interesting to finally decrypt a message that was sent during WWI that revealed the location of certain troops. However, that data would be of no military relevance today. So the argument exists not only about the relevance of data, but also how long the data is meaningful as compared to how long it might take to decipher. For these reasons, commercial encryption and even multiple layers of commercial encryption have been offered up as possible methods for protecting certain data types with minimal life expectancies. But performance and logical implementation has impeded these ideas. And so the debate will continue as to the classification of data in terms of lifecycle, importance of immediacy, and methods of transport. While policies will be diverse based on requirements, a structural data transport security framework is required as a baseline.

The Status Quo

Policies implemented on different data classes are simply recommendations penned to solve data security concerns. Implementation of these recommendations in a logical and sensible way is critical to ensuring that policies can be met in real world environments. In other words, policies do not address the problem in its entirety. For example, methods for sending government classified data types are proven solid in terms of encryption and key lifecycle management. Most however, do not necessarily address the vulnerabilities associated with traffic analysis and when they do, it is usually on a session by session basis. Providing Traffic Flow Security (TFS) on an ad hoc basis is in itself a vulnerability because if only highly secure network connections implement TFS, then highly secure connections can easily be identified. Full traffic flow security can only be obtained if the total network is obscured twenty-four hours a day, seven days a week regardless of the data classifications and policies.

A Complete Data Transport Strategy

Implementing strong commercial cryptography, consistency in performance, and complete obfuscation of network traffic at 98% bandwidth efficiency is at the core of a total data transport security strategy. Once this baseline for completely obscuring data in motion is implemented, all forms of traffic can pass unrestricted and with minimal impact to their current overhead constraints.

Classified Data

Traditional methods for sending classified data can be sent through the aforementioned secure infrastructure without compromising the integrity of the classified protocol. In fact, the additional layer of commercial cryptography combined with the obfuscation provided by TFS (e.g. consistent packet size and gap-filling, randomly generated traffic), provide additional data transport security on a full time basis.

Unclassified Data

As described earlier in this document, even the most benign communications can have security consequences. It should be understood that, especially for military applications, any data could result in a compromise. Masking unclassified data helps resolve issues related to data analysis and helps to mix traffic in a way that befuddles the listener. If data obfuscation was done only for sensitive data paths, then sensitive data lines could be easily determined by a listener. However, if all data is obfuscated, a listener cannot determine if the line is being used to send highly sensitive data or merely for the viewing of funny cat videos.

The Blurry Line Between Classified and Unclassified Data

This paper has presented a rationale for sending both Classified and Unclassified Data through a secure and efficient link. But much debate has arisen over the use of double, and even triple commercial encryption for sensitive data that simply cannot be distributed to "all soldiers" securely under current methods. Classified data transmission normally requires the use of classified devices and classified devices can only be handled by those individuals with the security clearance and the operational know-how. To compound the device usage issue, the misplacement or capture of a classified device can have major consequences to a mission. The limited usage of these devices is in direct conflict with the need to distribute huge amounts of data to a wide array of soldiers (with and without necessary security clearances). Using the commercial security baseline outlined in this paper, double encryption can easily be achieved. A simple encrypted VPN session from headquarters to soldier can be used to double encrypt these newly defined classifications of data if delivered through the secure baseline infrastructure.

Conclusion

The security of data transmission cannot be achieved by encryption alone. Completely securing the baseline infrastructure in an unobtrusive and unconstrained manner will pave the way for delivering and protecting all forms of data…classified, unclassified, and everything in between. At a minimum, the following four criteria must be met.

1. World class cryptography and key management
2. Protection of encrypted data as it moves through the infrastructure
3. Consistency in performance, regardless of frame size and data classification methods
4. Transparent use of bandwidth nearing 100% efficiency

Data transmission lines cannot be physically secured as they were during the day of dedicated physical links. Logical security methods must be implemented to ensure the same levels of security without the physical attributes that can be afforded by hard-wired links. Implementing the secure baseline outlined in this paper provides the means necessary to secure all data types from attacks as they move between physical and logical networks.