Transmission Security (TRANSEC)

iDirect

INTRODUCTION

We at iDirect recognize the critical need to protect the flow of communications to wherever the military and government agencies may operate. Wherever this may be, threat actors readily stand by to monitor, exploit or intercept communications for malicious intent. To mitigate this threat, we have been providing Transmission Security (TRANSEC) capabilities since the initial release of the Evolution software. With the release of 4.2, we have further enhanced our TRANSEC capabilities by extending protection to cover both one-way and two-way networks.

In combatant situations, where even a small "spike" in traffic can be a critical piece of intelligence, the need to mask any communications activity becomes apparent. The National Security Agency (NSA) has outlined the following vulnerabilities inherent in an IP-based TDMA transmission that must be addressed in order to provide true TRANSEC:

Channel Activity – The ability to secure transmission energy to conceal traffic volumes.
Control Channel Information – Disguise traffic volumes to secure traffic source and destination.
Hub and Remote Unit Validation – Ensure remote terminals connected to the network are authorized users.

What is TRANSEC?

TRANSEC prevents an adversary from exploiting information available in a communications channel without necessarily having defeated encryption.

TRANSEC requires all network control channels and Management & Control (M&C) data to be encrypted, and that any and all traffic engineering information be obfuscated from an adversary. For example, TRANSEC requires a communications channel to appear completely full to an adversary even if little or no actual data is flowing. This is contrasted with communications security (COMSEC); the actual communication (e.g. voice, video or data stream) is encrypted, but certain header information is sent in the clear. While the encryption is virtually impenetrable, the information in the IP header including the source address, destination address and, most importantly, the Type of Service (ToS) field are in the clear. With the IP header of an encrypted packet in the clear, an adversary can determine how much of the traffic stream is voice, video or data. More significantly, an adversary could determine when high-priority flash-override traffic has been initiated and from which location.

In a traditional SCPC (single channel per carrier) satellite network topology, achieving TRANSEC compliance is relatively straight forward. For SCPC connections, a bulk encryptor is employed to encrypt any data and control information traversing the network. The IP header of the packet would be encrypted by the bulk encryptor prior to being transmitted to the satellite. In addition, since an SCPC link is static, always on and no control information needs to be exchanged between the SCPC modems, all of the TRANSEC requirements are met.

In a TDMA network, TRANSEC compliance is more difficult. A TDMA network dynamically allocates bandwidth to remotes; therefore, there must be some type of control information transmitted to each device in the network. This control data containing traffic engineering information, as well as information available from an encrypted IP packet header, can be exploited by an adversary. For example, anomalous traffic volume to a specific remote can indicate new activity in that area while varying ratios of voice-to-data traffic can denote the distribution of intelligence (data) compared to lower priority voice traffic.

iDirect has implemented the following solutions in response to the security vulnerabilities of a TDMA Very Small Aperture Terminal (VSAT) network.

Masking Channel Activity

CHALLENGE

The first vulnerability that exists in a TDMA network is the availability of traffic engineering information. In an SCPC network, the link is static with no variation in transmission characteristics based on end user communications. An adversary looking at a satellite transponder with a spectrum analyzer will see a constant RF signal. This is contrasted with a TDMA network. A TDMA in-route carrier energizes and de-energizes as traffic flows and stops. The on-and-off nature of a TDMA in-route is the natural extension of the ability to allocate satellite transponder space to remotes that have transient demands. While this characteristic makes TDMA networks much more bandwidth efficient, it allows an adversary to determine peak periods of activity, identify unusual or unexpected activity spikes, and identify locations of remotes that have remained quiet for a period of time and suddenly experience increased traffic volumes. The obvious risk in having this information in the hands of an adversary is the potential to extrapolate timing, location and scale of a strategic activity.

SOLUTION

We at iDirect have implemented free slot allocation in our TDMA bandwidth distribution algorithm. With free slot allocation, an adversary snooping for satellite transponder energies will see a constant "wall of data" regardless of traffic profiles. As the name implies, free slot allocation keeps the in-routes active regardless of actual traffic flows. Free slot allocation preserves the efficiencies of a TDMA system while obfuscating actual traffic volumes, negating the risk of using transmission activity as an intelligence gathering mechanism.

Obfuscating Acquisition Activity

CHALLENGE

The rate at which remotes acquire into a network can provide critical information to an adversary about troop activities. All TDMA networks provide a dedicated channel for remote acquisition activity. If adversaries monitor the activity in this channel, they will be alerted to troop movements by a flurry of acquisition activity.

SOLUTION

We exceed TRANSEC requirements by addressing acquisition activity vulnerability. The iDirect acquisition algorithm inserts dummy bursts from remotes already in the network and intentionally skips acquisition bursts at times of high activity, ensuring an adversary sees only a random distribution of acquisition activity. The iDirect acquisition algorithm goes a step further by randomly varying the dummy burst's frequency, timing and power. This randomization makes sure an adversary cannot distinguish between a dummy burst and actual acquisition activity.

Control Channel Information

CHALLENGE

A great deal of traffic volume and priority information can be gleaned by examining the in-band or out-of-band control information within an encrypted Time Division Multiple Access (TDMA) network. As previously discussed, the IP header of a packet contains source, destination and priority information. In order for a TDMA network to provide the quality of service (QoS) needed to support real-time traffic, data quantities and prioritization information must be gathered. This information could be more useful to an adversary than channel activity data because it is specific enough to delineate between general communications like email and web traffic, versus tactical communications like voice and video.

SOLUTION

The only solution for this vulnerability is to completely encrypt all Layer 2 information as well as any control information disseminated to the remotes. The encryption methodology must be secure enough to thwart an adversary long enough that the data becomes old and unusable. We have implemented Federal Information Processing Standard (FIPS) 140-2 certified 256-bit keyed Advanced Encryption Standard (AES) for all Layer 2 and control information. The encryption of the Layer 2 frames has a side benefit of re-encrypting the data payload. Therefore, the transmitted IP header itself is AES-encrypted. Additionally, the iDirect TRANSEC TDMA slot is a fixed size, again to obfuscate any traffic characteristics. This Layer 2 encryption solution solves all existing control channel vulnerabilities. The iDirect Layer 2 encryption method goes a step beyond to feature over-the-air (OTA) key updates and a unique Layer 2 frame format, including an Initialization Vector that ensures randomization of repetitive data streams. The net result is that adversaries are precluded from detecting any repetitive pattern, which can aid in deciphering encryption algorithms.

Hub and Remote Authentication

CHALLENGE

Another vulnerability of a TDMA VSAT system is the concept of Hub and Remote validation. In traditional SCPC architectures, a link remains active for very long periods of time when it is established. Because these connections are fixed, and there is a significant level of coordination between personnel commissioning the SCPC, a high degree of confidence exists that an adversary is not trying to assume the identity of a trusted entity. In TDMA networks, remotes are routinely coming into and dropping out of the network. This is especially true of networks with mobile or itinerant terminals where terminals are located in moving vehicles, aircraft or maritime vessels. This type of dynamic environment gives an adversary a greater opportunity to obtain a VSAT remote through licit or illicit channels, spoof the device ID and insert a rogue remote into a secure network. Equally feasible is an adversary acquiring a VSAT hub terminal and coaxing a blue force remote into the adversary's network.

SOLUTION

To mitigate this risk, we have implemented X.509 digital certificates on TRANSEC remotes. An X.509 certificate utilizes RSA public key cryptosystem. With this cryptosystem, two related keys are generated: one private key and one public key. The functionality of these keys is so that anything encrypted with the public key can only be decrypted with the private key, and anything encrypted with the private key can only be decrypted with the public key. In the iDirect system, X.509 certificates can be generated via the NMS server. Certificates are placed on all TRANSEC line cards and Protocol Processors as well as on the remotes. The hub system keeps the public keys of each remote configured to operate on the hub, and the remotes have the public keys of each hub device. During network acquisition, the remote encrypts its X.509 certificate with its private key, and the hub verifies by decrypting the certificate with the remote's public key and vice versa. This process ensures a remote is not only authorized to operate in the network, but that the hub is a trusted entity.

Operational Implementation

CHALLENGE

Implementing security and ensuring all security policies are followed can be a burden to the soldier in the field. Implementing TRANSEC and performing key

management is no exception. Challenges one would face in operating a TRANSEC network include creation, distribution and revocation of X.509 certificates; ACQ and Data Channel key generation, distribution and management; and zeroizing modems. A robust TRANSEC network also requires the use of at least two network-wide keys: The ACC Key for acquisition, and the DCC Key for the data channel. A long-lived, user-generated passphrase is used to protect the keys during initial commissioning. The use of front panel displays to enter the passphrase and external key fill mechanisms places an undue burden on the warfighter and introduces security vulnerabilities.

SOLUTION

We have implemented a FIPS-approved software method of key generation and automatic, OTA key distribution protocol. Not only does the software-based key generation and key distribution mechanism make TRANSEC operation simpler and more convenient for the warfighter, it makes the system much more secure by removing a human from key distribution.

Another advantage of automatic key generation and distribution is that it seamlessly enables a global communications-on-the-move (COTM) TRANSEC network. By automatically generating and distributing new acquisition passphrases, a single, dynamic passphrase can be utilized across global networks.

Additional Security Measures

FIPS 140-2

The FIPS 140-2 is a U.S. government security standard for accrediting cryptographic modules. The standard is published by the National Institute of Standards and Technology (NIST).

FIPS 140-2 provides stringent third-party assurance of security claims on any product containing cryptography that may be used by a government agency. FIPS 140-2 establishes the Cryptographic Module Validation Program (CMVP) as a joint effort between NIST and Canada's Communications Security Establishment (CSE).

FIPS 140-2 specifies four levels of security when it comes to the design and implementation of cryptographic modules. As described by NIST, the following is a high-level overview of these security levels:

Security Level 1 is the basic level of security. No specific physical security features are required, and only one approved security function algorithm is required.

Security Level 2 requires tamper-evident coatings or seals that must be broken to gain access to the cryptographic keys and critical security parameters.

Security Level 3, in addition to the requirements of Level 2, physical security mechanisms are required to be able to detect and respond to attempts at physical access or modification of the cryptographic module.

Security Level 4 requires a complete envelope of protection around the cryptographic module with the ability to detect and respond to all unauthorized attempts at access.

In addition to the hardware requirements described above, FIPS validation applies to the cryptographic solution as a whole, including the operating system and software.

Enhancing TRANSEC and Security with the 9-Series and DLCs

With the release of our new 9-Series Satellite Routers and Defense Line Cards, we have expanded our existing FIPS 140-2 certification from Level 2 to Level 3 from our previous line of products. As part of the effort, we developed a TRANSEC module designed to meet the stringent FIPS 140-2 Level 3 requirements as defined by NIST. Through hardware and software development, the embedded, and yet independent, TRANSEC module on the 9-Series and DLCs operates through a separate and trusted path from all other interfaces on the product. The module also features a strong physical security measure for tamper prevention and the capability to zeroize the security keys or critical security parameters (CSPs) stored on the module itself. If required, the revocation or zeroization of the keys can be accomplished either OTA by the hub operator or locally on the remote by authorized personnel.

One-Way Networks

We have further enhanced our TRANSEC capabilities by securing one-way broadcast transmissions. Based on their encapsulation method, LEGS, the iDirect platform can provide the same level of security for one-way networks to that of two-way networks as described earlier. The 900 and 9350, with its dual-demodulator support, are capable of dual-domain TRANSEC; the ability to establish two independent chains of trust (sets of X.509s) between two different CAs. An example use case of this feature would be one demodulator on a two-way TRANSEC network while the second demodulator receives a separate one-way TRANSEC-secured broadcast. Elliptical Curve Cryptography (ECC) is used for key generation along with X.509 certificates for authentication in each security domain.

iDIRECT

13861 Sunrise Valley Drive, Suite 300, Herndon, VA 20171
+1 703.648.8000 | +1 866.345.0983 | www.idirect.net