2015 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
VEHICLE ELECTRONICS AND ARCHITECTURE (VEA) TECHNICAL SESSION
AUGUST 4-6, 2015 – NOVI, MICHIGAN

SUBSCRIBER CONDITIONAL ACCESS SOLUTION FOR DYNAMIC IN-FIELD PKI AUTHENTICATION

Mr. David Jedynak
Chief Technology Officer, COTS Solutions
Curtiss-Wright Defense Solutions
Austin, TX

ABSTRACT

This paper will lay out the critical challenges of in-field Public Key Infrastructure (PKI), namely Integrity, Availability, and Confidentiality, and will assess multiple conceptual solutions against them. The history and mechanisms of Subscriber Conditional Access will be detailed to provide understanding of this technology. Mapping of PKI data into a Subscriber Conditional Access system will be provided, showing a solution which meets all challenges. Analysis of organizational hierarchies, dynamic control latency, and required data bandwidths will be provided. Finally, a reference architecture showing how to implement a Subscriber Conditional Access system for Dynamic In-field PKI Authentication will be provided.

INTRODUCTION

Securing information systems using Public-Key Infrastructure (PKI) is a well understood and widely practiced technique for authenticating users and processes for communication and access. This includes secure authentication and communication over untrusted public networks, such as the internet. The power of PKI is that it uses a trusted Certificate Authority (CA) to validate both identity and authorization for users and processes attempting to access or communicate with various systems. Administration of the CA is performed under controlled processes with trusted personnel, and provides such flexibility as changes to a given identified certificate's authorization levels as well as complete revocation of all privileges in a dynamic manner, always ensuring that the current result of an authentication request matches the current desired state of authentication for the particular end-user or processes. In addition, the central repository of user and process level access means that concepts such as Single Sign-On (SSO) and mobile profiles for access which follow the user anywhere (including the user's own devices) rather than the local hardware are possible. The use of smart cards (e.g. Common Access Card) for two-factor authentication relies on and takes advantage of the power of PKI.

Although this infrastructure works very well in an "always-on / always-connected" environment such as a fixed installation with persistent and controlled data-links (e.g. physical lines) to the Certificate Authority, the use of PKI in-field for authentication has numerous challenges, as shown in Figure 1. In an environment where a persistent connection (Availability) to the central CA is not possible, not desired, or susceptible to denial, real-time authentication to the CA doesn't work. Local replication, thus exposing the entire CA to compromise, is not desirable. Without the power of PKI, other far less secure methods are used to secure systems for in-field use. The current practice of fixed role-based authentication with essentially semi-permanent passwords unfortunately exposes systems to compromise, and makes secure key management for various levels of classification basically impossible.

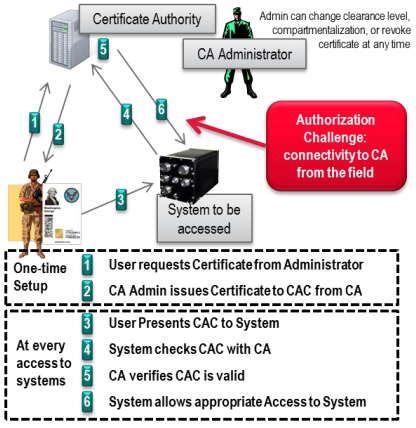

Figure 1: DoD Public Key Infrastructure with CAC highlighting in-field PKI challenge

Fortunately, the direct subscriber broadcast industry (radio, television) has implemented and matured over 30 years a mechanism of Subscriber Conditional Access to dynamically distribute protected content from a central head-end to a geographically diverse set of subscribers which are all dynamically managed from the head-end, as first described in [1]. This paper shows how the system of Subscriber Conditional Access can be used to dynamically replicate PKI Validation Authorities to assets in-field. This solution keeps the CA secure, ensures availability of the authentication solution, and protects the user and process level keys from compromise if an asset is lost in-field.

CONCEPTUAL SOLUTIONS AND PROBLEMS

The key problem is the inability to interface with a trusted authority, in the field, to provide authentication. This problem can be broken down into three standard security constructs: Integrity, Availability, and Confidentiality.

A key reason to use PKI is that it uses a mutually trusted 3rd party to both issue and validate security artifacts. It's absolutely critical that the Certificate Authority (and a subordinate Validation Authority) maintain Integrity, both at a given instant and across time. For example, the CA must be under trusted control and provide trusted information regarding various user credentials, but it also means that the CA must be under trusted control at a future state as well, representing future changes to the security artifacts based on changes to the status of users or information (e.g. Revocation of permissions). The CA must have a mechanism to change state and provide visibility to that state change. Normally, the CA is administered locally by trusted administrators to affect state changes, and since the CA is always available to users, changes are immediately visible to the user-base. In the field, it's more difficult to maintain this Integrity, as the state will become stale with any changes. Any solution providing authentication in the field needs to address this.

Worse than Integrity problem is the Availability problem. Really, the in-field issue is one of availability – the CA is just not available for connection, which will result in degraded functionality for users, or a complete lack of functionality. Any way of providing authentication in the field needs to solution needs to address the availability issue, while also considering the deliberate attempt to deny access in order to disrupt the PKI system.

The Confidentiality of the CA is itself a critical concern. Since the CA contains a large amount of user data and private keys, and is in charge of both issuing and validating credentials, it needs to be protected. If a CA falls under the control of an unauthorized force, the entire PKI infrastructure can be subverted. At a more localized scale, a solution for in-field PKI needs to recognize its fundamental reason for use – to authenticate users and processes as they attempt to access other users and processes. The confidentiality of those users and processes needs to be ensured through the PKI solution, just as if the CA was always connected (e.g. a user's authentication is renewed / revoked based on changes to state or information). Any solution to address the in-field use of PKI needs protection from compromise for both the PKI system and the users and processes it serves.

These three dimensions are really best understood with a set of conceptual solutions, however imperfect, as shown in Table 1.

Table 1: Conceptual Solutions and Issues for In-Field PKI

Solution	Integrity	Availability	Confidentiality
Require CA link at System Boot, then continue at that authentication level	No means for revocation / change	CA Link jamming / availability prevents system uses	Captured System is authenticated - users and processes are wide open to compromise
Mirror CA Locally	No means for revocation / change	Available but requires large storage for CA	Entire CA Compromised with System
Require CA link at all times	Full control of CA at all times by CA admin	CA Link jamming / availability prevents system uses	CA protected, but link is a threat vector
Ideal solution for In Field PKI CA Solution	Full Control of CA at all times by CA Admin	Available without needing large local storage resources	CA protected, threat vectors minimized

The first conceptual solution – requiring a link to authenticate all platform users and systems at startup / boot – assumes that the platform will start up in an environment where a link can be established. This is problematic, since a platform will not always be "behind the wire" when starting up. In addition, simply jamming this live connection could render a platform unable to start up.

The other problems with this solution are that the platform authenticates at a single point in time, and then essentially runs free from the rest of the authentication infrastructure. For instance, a subsystem on the platform accesses Secret or Top Secret / Compartmentalized information resident on the platform. This subsystem (and user) authenticates via PKI and the platform starts its mission. If this platform is then captured without another start-up cycle, then the classified information will still be wide open. In this approach there is no way to "turn off" the access for this user via revocation of security credentials.

The second conceptual solution – mirroring the CA locally – provides a solution to the availability problem, but still

exacerbates the integrity and confidentiality problems. Having a local CA means validation of identity and status can be handled at any time (assuming the CA is functional), but it means that it is immediately stale as soon as it is mirrored from the root CA. No administration can be performed on it in the field, resulting in a same outcome for the classified information above. In addition, the CA itself is now open to compromise, as it sits on the platform, and is open to interrogation and manipulation by adversaries who gain control of the platform. In addition, mirroring the entire CA on the platform can represent a significant storage challenge.

The third conceptual solution – requiring a link to the CA at all times – is really the core problem, but is included as it is illustrative of partial solutions and further complications. A platform which requires a link to the CA at all times to continue operations (or at least operations which require authentication for any or non-degraded operation) provides a solution for the integrity issue, but makes the availability situation untenable. In addition, PKI infrastructure itself is opened to compromise since the very nature of the query / result link between the platform and the CA means the CA can be interrogated and probed for compromise from a captured platform.

These different conceptual solutions illustrate a set of requirements for a dynamic in-field PKI solution:

1) Full control of the CA at all times by the CA Administrator
2) On platform (local) availability with a reasonable requirement for storage
3) The CA itself is protected, with minimized threat vectors and a subset of records on the local platform.

SUBSCRIBER CONDITIONAL ACCESS

In the 1970s, a new business model emerged in the broadcast industry – the delivery of premium subscriber-based content over a broadcast channel (e.g. UHF or a common cable television system). Generating direct revenue from subscribers on broadcast content required a way of limiting the reception and viewing of that content only to authorized subscribers. Given that the signals were broadcast over the air or carried on a common CATV, a technological method was needed to only allow subscribers access to the signals. The initial solutions focused on scrambling of signals in such a way that only the appropriate equipment (a "descrambler") would allow a viewer to see and hear the content. Subscribers would be issued the appropriate descrambling equipment, and upon returning it, their subscription charges would cease.

Although this system worked, it had a significant vulnerability – unauthorized descrambling equipment, built via reverse engineering techniques and sold by 3rd parties as a way to receive and descramble the content for free. Regardless of the legality, this black box piracy market resulted in the need for a more dynamic solution.

In 1982, Tony Wechselberger and a team led by Dr. Leo Jedynak of Oak Industries filed a patent for a "Multi-layer encryption system for the broadcast of encrypted information" [1]. The core of this far reaching and still referenced patent addressed the challenge of the broadcast delivery of information to subscribers, introducing the concept of subscriber conditional access. This approach will be familiar to anyone using satellite or cable broadcast systems, in which channels or time blocks of content (pay-per-view) can be turned on or off remotely, based on subscriber fees and payments. The patent was broad in its description of applications to include "other types of digital or digitized data such as computer software, games, radio programs, computer data bases, etc." [1]

This patent laid out two fundamental concepts:

- Broadcast content can be encrypted with global keys
- Global decryption keys can be individually distributed

Essentially the same as analog scrambling and the black box descramblers, the encryption of broadcast content could be defeated with sufficient reverse engineering and resources. The critical development was the additional concept of changing the encryption key as needed, then securely distributing the new global key to individual subscribers as a method of defeating a compromise. This relies on the fact that each individual receiver box has its own unique identity with a set of unique encryption keys.

The overall system has a defined content channel and control channel. The content (broadcast) channel is used for globally encrypted content, while the control channel is used for individual communication with each receiver box. The control channel is intended to be multiplexed within the broadcast itself, but could be out-of-band. Global keys (and other receiver commands) are uniquely encrypted and distributed via the control channel, providing the fundamental infrastructure for subscriber conditional access, as shown in Figure 2.

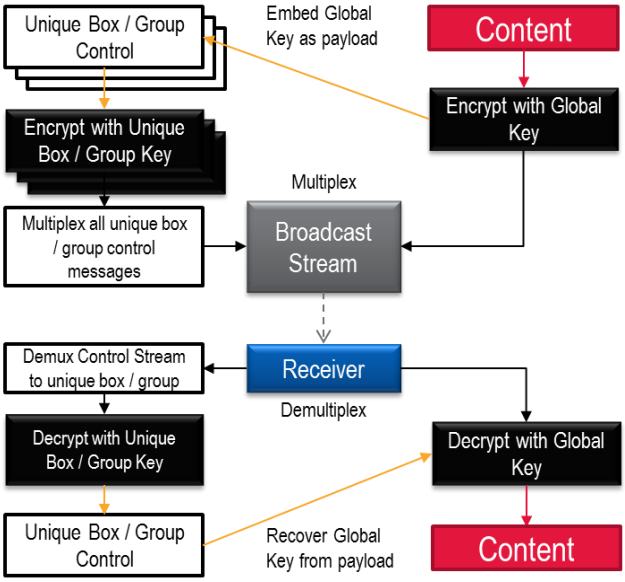

Figure 2: Basic operation of Subscriber Conditional Access System

An individual box (or group) can be authorized or de-authorized simply by providing it or not providing it the current global decryption key. In the case of piracy, the global key can be changed, and the new key simply not delivered to the pirate box (most likely a clone of an authorized box). The important concept here is that it is not necessary to actually de-authorize a box remotely, but to simply not provide it with updated information.

This foundational scheme, the basis of all modern subscriber conditional access systems, including those over modern networks, is directly applicable to the challenge of dynamic in-field PKI authentication.

APPICATION TO PKI

Although the basis of Subscriber Conditional Access was in the broadcast of dynamic audio / video content, the patent was clear to note that it be used to deliver other content as well. The key challenge to in-field PKI authentication is the placing a trusted authority into the field. Applying subscriber conditional access to PKI, the read-only version of a Certificate Authority, the Validation Authority (VA) is the dynamic content, as shown below in Figure 3.

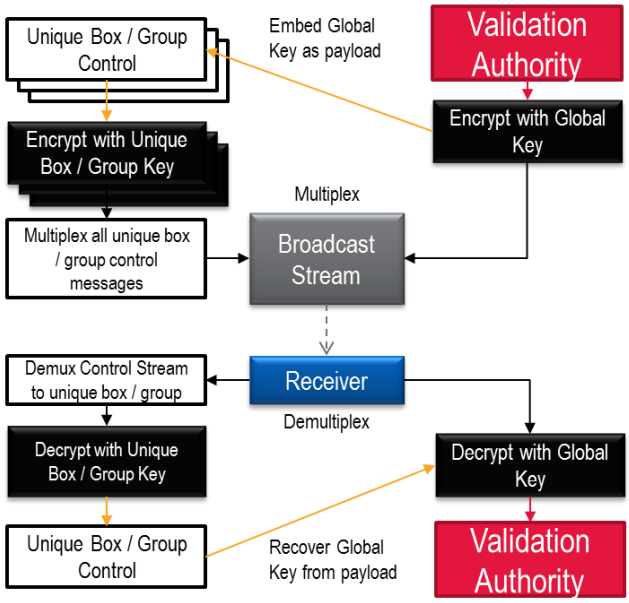

Figure 3: Application of Subscriber Conditional Access to PKI Validation Authority Broadcast

The application of Subscriber Conditional Access methods addresses the challenges of Integrity, Availability, and Confidentiality.

By dynamically distributing the VA, which is essentially a snapshot of the current state of the Certificate Authority, state changes (including revocation) are propagated to the platforms in-field. Since a snapshot of the VA is onboard a platform, availability can be addressed through various aging properties. Since the VA is a broadcast of a snapshot of the CA, no back channel exists to interrogate or probe the CA, and the decrypted broadcast VA itself can be locally encrypted with the appropriate Data-At-Rest solutions for the platform, both of which address the main concerns of confidentially. The application is shown in Figure 4.

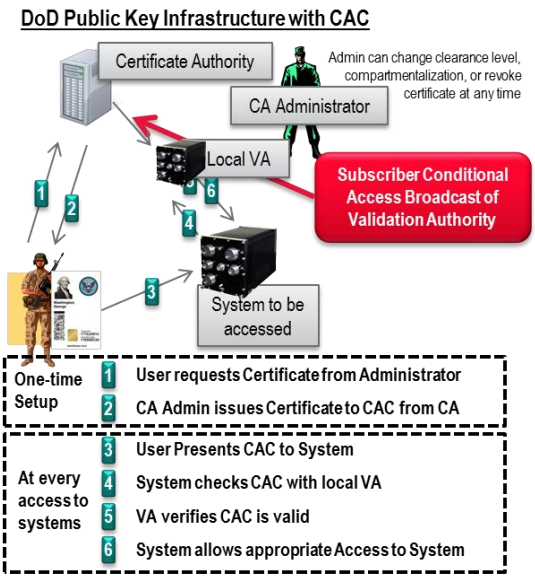

Figure 4: Application of Subscriber Conditional Access for Dynamic In-Field PKI Authentication

Optimization and performance concerns are discussed in the following section.

OPTIMIZATION AND PERFORMANCE

There are handful of critical optimizations and performance parameters involved in the application of Subscriber Conditional Access to Dynamic In-Field PKI. Fundamentally, all of these center on the amount of information which is broadcast versus the channel capacity. If an infinitely sized broadcast channel were available, the entire VA could be distributed to all platforms instantaneously, providing instant control over authentication changes; however, this is not the case. The channels are limited, so the update rate and the subset of the VA must be tuned for a balance of performance.

Three main parameters must be considered:

- The size of the Validation Authority (content), which is just a large database of PKI certificates
- The size of the control message stream (control), which is an aggregate of control messages, PKI keys, and subscriber unique identifiers
- The update rate

The overall bandwidth required is determined as follows:

$$B = N \frac{(C + K + I + F)}{U}$$

Where:

$$B = Bandwidth$$
$$N = Number\ of\ Subscribers$$
$$C = Certificate\ Size$$
$$K = PKI\ Key\ Size$$
$$I = Unique\ ID\ Size$$
$$F = Control\ Flags\ Size$$
$$U = Update\ Period$$

Implicit in this is the number of subscribers (N), which affects both the Validation Authority size and the control message stream size. For simplicity, it is assumed that the number of subscribers is the same or greater than the number of records in the VA. This means that all subscribers are assumed to be potential users attempting to authenticate to the common Validation Authority. This provides an upper bound on the bandwidths required.

The Certificate Size (C) is the actual size of a PKI certificate containing the public key of an individual subscriber. An organized set of these comprises a Validation Authority. The content stream size is the product of the number of subscribers and the size of the certificates.

$$Content\ Stream = N \times C$$

The size of the current global PKI Key (K) key is included because it is payload within the control stream. Each subscriber has a unique identifier (I) embedded in the control stream, enabling a subscriber to recognize and extract its specific control message from the multiplexed streams. Additional control information, such as an aging parameter, is encoded as flags in the control message (F). The control stream size is the product of the number of subscribers and the sum of the overall control channel elements:

$$Control\ Message\ Size = K + I + F$$

$$Control\ Stream = N\ (K + I + F)$$

The aggregate of the content stream and the control stream is the total stream:

$$Total\ Stream = (N \times C) + N(K + I + F)$$

$$= N(C + K + I + F)$$

The update period (U) is the factor that determines how dynamic the Validation Authority is in the field. Slower update periods mean a longer time that a stale validation authority exists in the field. The directly impacts how long it takes for an authorization change (granting access or revoking it) will take to propagate. The bandwidth required is inversely proportional to the update period (or alternatively the bandwidth required is proportional to the update frequency).

The number of subscribers is a very critical parameter to the usability of the system. There are a number of optimization goals, some of them conflicting:

- Minimize the size of a sub-Validation Authority for broadcast via communication networks
- Minimize the number of unique subscribers to be addressed
- Maximize scope (size) of sub-VA for fluidity of personnel across assets in theater operations
- Minimize longevity of sub-VA validity to reduce compromise windows
- Maximize longevity of sub-VA validity to reduce out-of-communication (LOS / Jam) availability
- Optimize to communication channel bandwidth as well

No matter how the optimization is done, the fundamental structure remains the same: broadcast a subset of the overall CA in the form of a subordinate VA with the appropriate size of subscribers (records).

The optimization reduces down to the question of where to cut the organizations and how often to update, as shown in Figure 5.

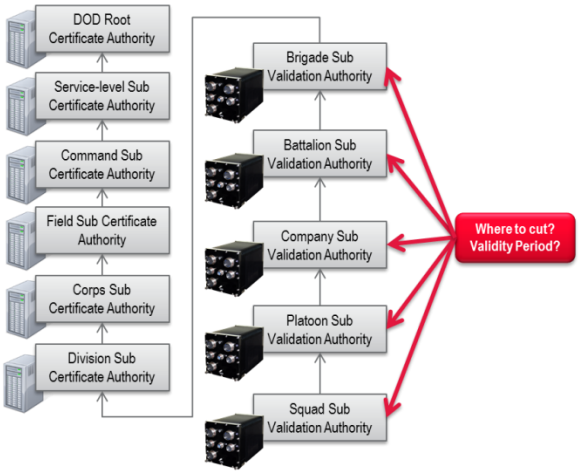

Figure 5: Subordinate VA hierarchy

Balancing fluidity of personnel and assets with overall bandwidth, it is assumed that the sub-VA will be at the Battalion level. The assumption is made that the total number of subscribers and certificates is approximately 1000 (assume personnel and platforms are both viewed as subscribers, and all have unique certificates in the validation authority). In addition, this provides for a number of certificates for cross-organizational interfaces.

Table 2 presents additional assumptions for a calculation of bandwidth required:

Table 2: Bandwidth Calculation Assumptions and Calculated parameters

Parameter	Value	Notes
Number of Subscribers (N)	1000	Battalion
Certificate Size (C)	2000 Bytes	No compression
Content Stream Size	*2 Megabytes / 16 Megabits*	$N \times C$
Global PKI Key Size (K)	2048 Bits	Largest in general use
ID Size (I)	64 Bits	Long Integer
Control Flags Size (F)	64 Bits	Long Integer
Control Message Size	*2176 Bits*	$K + I + F$
Control Stream Size	*2.176 Megabits*	$N(K + I + F)$
Total Stream Size	*19 Megabits*	$N(C + K + I + F)$ *Rounded up from 18.176*

The important result in this table is the Total Stream Size for a Battalion = 19 Megabits. This means that to distribute the globally encrypted Validation Authority to 1000 subscribers using a Subscriber Conditional Access mechanism requires 19 Megabits. Note that this also addresses the earlier data storage concern of mirroring a Validation Authority locally – it's less than 3 Megabytes for the current image.

The tuning of the update period is the next step. Table 3 presents calculation of the required bandwidth to broadcast the Validation Authority with Control Channel at the given update period (U):

Table 3: Overall Bandwidth Required to Deliver 19 Megabit VA & Control for Various Update Periods

Update Period (min)	Update Period (sec)	Bandwidth Required (Kilobits / second)
1	60	316.7
10	600	31.7
60 (1 hour)	3600	5.3
720 (12 hours)	43200	0.44
1440 (1 day)	86400	0.22

Taking the case of a 1 hour update period, meaning the entire Battalion will be updated with a refreshed snapshot Validation Authority over the course of one hour, the total bandwidth required to travel over any available waveforms is a sustained 5.3 Kbps. This is a small trickle compared to the capabilities of CDL and WIN-T, and can be supported by Link-16, and could be supported by a dedicated SINCGARS.

All of this assumed that the number of subscribers and the number of certificates were equivalent. If that was not the case, and the number of subscribers was actually 10% the number of certificates, the end result is not significantly different. The overall control channel would be approximately 220 kilobits rather than 2.2 Megabits. The change to the Total Stream size would not be significant, reducing to ~16.2 Megabits from an original 18.2 Megabits. The 1 hour update period required bandwidth would correspondingly reduce to 4.5 kbps from 5.3 Kbps.

Interestingly and somewhat ironically given the origin of subscriber conditional access, the current Advanced Television Standards Committee (ATSC) transmission rate for digital television broadcasts provides a payload of 19.2 Megabits / second. This means that commercially available digital television transmitters and tuners could be used as a method to distribute a complete 19 Megabit Validation Authority + Control snapshot once a second if the video payload is replaced with the VA + Control payload.

REFERENCE ARCHITECTURE AND USAGE

Implementing the Subscriber Conditional Access Solution for Dynamic In-Field PKI Authentication requires addition of software / hardware provisions both in the field, and within the current Certificate Authority infrastructure.

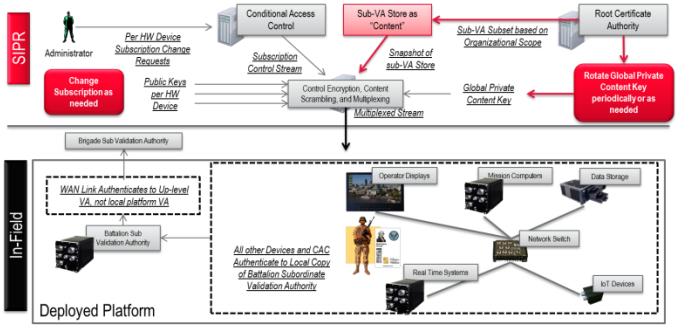

Figure 6: Reference Architecture

The Reference Architecture is intended a baseline for discussion. Returning to the example of classified information, the administrator of the CA will determine (or be commanded) that a particular platform in the field has been compromised, and it should be de-authorized from accessing that information. The CA is updated to revoke the certificate of that platform (or the onboard subsystem). It can no longer authenticate to the other system providing the classified information since it will no longer have appropriate credentials in the local sub-VA. This revocation will replicate to the local sub-VA during the standard update rate.

On the other hand, and adversary may realize that the local VA is being updated remotely via the Subscriber Conditional Access system and will attempt to block future updates, attempting to freeze the system and local sub-VA in a current state (e.g. to allow the continued use of certain advanced systems which rely on PKI based cryptographic ignition). This is where an aging property, previously referenced, is of critical use. There are three general settings for an aging property in this context:

1) After an update period without update, discard the local VA
2) After an update period without update, hold the local VA indefinitely.
3) After an update period without update, hold the local VA for a predefined time

This is where the control message flags become useful. The aging property and time-out behaviors can be set remotely be the administrators. For example, if a platform or personnel are expected to go on a mission without radio contact for no more than 3 days, an aging property can be set for "hold for 3 days" (or more for margin). During this window, if there is any compromise, the window will still be open, but it will close at the end of the aging period.

It is assumed that the reference timing (clock) of the local sub-VA is adequately protected, and that contingencies are provided (e.g. anti-rollback). This topic and others related to the integrity of the local sub-VA hardware and software is beyond the scope and intended audience of this paper.

The use of the Common Access Card is for illustration purposes of two-factor authentication and PKI. Its use a security token in-field is problematic given its other roles for identification. Other identification schemes, such as biometric / RFID FOB can be used in a similar fashion, replacing the various identifiers in the Validation Authority as necessary.

CONCLUSION

Subscriber Conditional Access is an ideal construct for distributing and managing authentication credentials in-field.

The 30+ year maturity of this approach and continual usage across multiple applications makes it a low-risk COTS approach to solving isssues of in-field cybersecurity and authentication. The bandwidth required for delivery to the field is minimal compared to multiple communication system, and does not require the development of new communication channels – in fact , it can leverage commercially available communications channels as well.

The use of the Subscriber Conditional Access Solution for Dynamic In-Field PKI Authentication addresses the Integrity of a PKI Solution through regular broadcast updates, ensures a measure of availability based on update timing and aging parameters as well as using established communication channels, and ensures the continued confidentiality and isolation of the root Certificate Authority from in-field risks.

REFERENCES

[1] A. Wechselberger, L. Bluestein, L. Jedynak, D. Drake, and L. Simpson, "Multi-layer encryption system for the broadcast of encrypted information," U.S. Patent 4531020, issued date July 23, 1985.