Secure Wireless Key Management for MAC-Layer Security and First Responder Credentialing[1]

Tim Draelos, Mark Torgerson, Pete Sholander
Sandia National Laboratories
Albuquerque, New Mexico
{tjdrael, mdtorge, peshola}@sandia.gov

Soumitri Kolavennu, Denis Foo Kune
Minneapolis, Minnesota
{soumitri.kolavennu, denis.fookune}@honeywell.com

Abstract: Critical infrastructure utilities are beginning to use wireless mesh networks in their monitoring, and sometimes control, networks. Establishing secure, distinct communication links between nodes within wireless mesh networks is important to avoid compromise of the entire network when a single node is compromised. In addition, first responders must be allowed authorized access to the network in a way that maintains security during and after emergencies. This paper describes key management protocols to support 1) distinct medium access control (MAC) layer link encryption and authentication and 2) time-limited credentialing of first responders within wireless mesh networks.

Secure communication in the wireless mesh network is dependent upon the secrecy and strength of the keys used in cryptographic algorithms for confidentiality and integrity of data and for authentication of authorized users. Trust in network security is rooted in a central trust authority (CTA) that is responsible for a) verifying the credentials of authorized wireless mesh nodes, b) generating and maintaining certificates for mesh nodes, and c) generating and maintaining time-limited credentials for first responders. Leveraging trust in the CTA, we use public key techniques to exchange a unique symmetric key between network neighbors. For first responder credentialing, we take advantage of the Remote Authentication Dial-In User Service (RADIUS) protocol to terminate sessions after a specified time has elapsed. The goal of this work was to provide another layer of communication security in mesh networks and to enable a new capability for first responders by leveraging existing cryptographic protocols.

Keywords: Wireless Communication Security, Encryption, Authentication, Key Exchange, RADIUS, First Responder, MAC-layer

[1] This work was supported under Award Number 2003-TK-TX-0003 from the U.S. Department of Homeland Security, Science and Technology Directorate. Points of view in this document are those of the authors and do not necessarily represent the official position of the U.S. Dept. of Homeland Security or the Science and Technology Directorate.

1 Introduction

This paper presents results of a partnership between Sandia National Laboratories and Honeywell International, namely enhancements to improve the reliability and security of existing wireless mesh network systems used for process monitoring and control in critical infrastructure facilities, see Figure 1. Improved security is achieved through the use of authentication and secure key exchange protocols between the wireless routing nodes. Enhanced availability and reliability is achieved via a unique non-overlapping, redundant, mesh-routing protocol for wireless networks. The resulting network is well-suited for robust and secure industrial applications as well as supporting emergency responders in the event of a crisis situation at a facility.

The focus of the current work is on a key exchange protocol to support Medium Access Control (MAC)-layer communication security and on time-limited credentialing of first responders. Establishing secure distinct communication links between routing nodes in a mesh network is important to avoid compromise of the entire network when a single node is compromised. A secure, reliable wireless mesh network will allow plant workers at critical infrastructure facilities to access the wireless network, increasing productivity and accuracy by keeping them connected to the plant's information system. During emergencies, first responders will also be allowed secure authorized access to the network for a limited amount of time in order to maintain situation awareness (SA) in support of emergency cleanup, rescue, repair, safety, medical care and recovery.

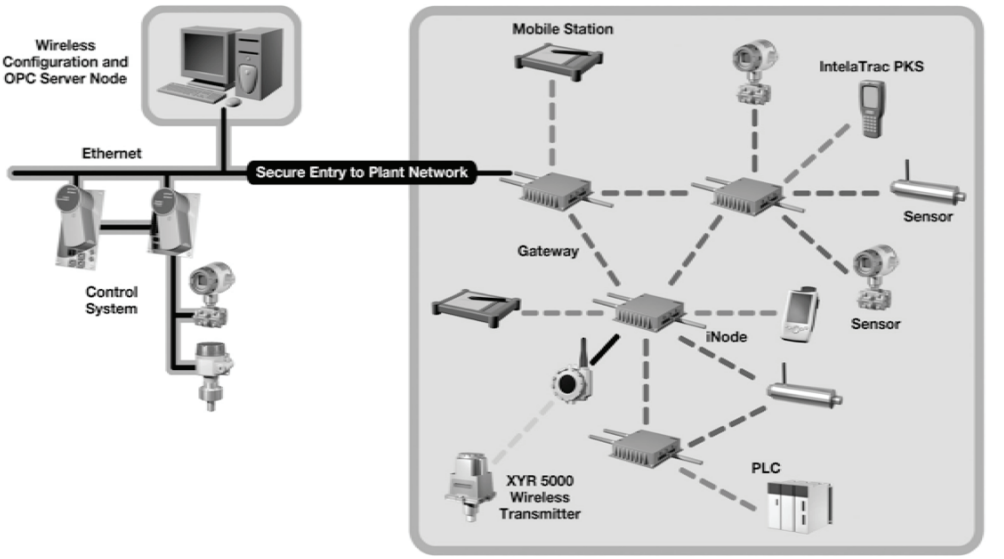

Figure 1 – Multifunctional Architecture for Industrial Wireless Sensor Network

1.1 Organization of Paper

In Section 2, we present the need for security in wireless mesh networks, including a matrix that compares the current state-of-the-art and our presented security/reliability enhancements versus threats against, and degradations found in, process control systems. Section 3 contains options for key management of networks and the security ramifications of each. In Section 4, the protocol for MAC-layer security in wireless mesh networks is presented and in Section 5, the protocol for time-limited first responder credentialing is presented.

2 Wireless Mess Networks

Wireless mesh networks are a collection of wireless, e.g., radio, communication devices organized in a mesh topology. A mesh network is a network that provides continuous, redundant connections between nodes via multiple hops or routing from source to destination. This allows communication to continue and the network to stay fully connected despite the presence of broken nodes and blocked paths. The type of network of interest here is an "infrastructure wireless mesh network", where nodes operate as routers to support communication for attached clients such as between sensors, controllers, and first responders.

The particular network used in the Sandia/Honeywell partnership is OneWireless™, a self-forming, self-healing, structured wireless, mesh-sensor network designed for use in harsh industrial environments. Developed by Honeywell on a cost-shared program with the U.S. Department of Energy, the objective was to develop a secure wireless network with the reliability of a wired network that could be used to economically reduce energy consumption and pollution, while simultaneously increasing the throughput of energy-intensive manufacturing processes. The OneWireless system offers the potential for increased network reliability and security, while also providing a potential information and communication conduit in which pertinent SA data can be seamlessly delivered to first responders in emergency situations.

The OneWireless mesh network is composed of a central mesh of routing nodes (iNodes) as shown in Figure 1, each possessing at least two radios. One radio is used to talk on the mesh network. The other radio is used to allow first responders with 802.11b/g devices (such as laptops or PDAs) to connect to the mesh. An iNode can use its access point MAC address and its mesh radio MAC address as credentials for authentication purposes.

The overall OneWireless system architecture is a multifunctional network capable of supporting both 802.15.4-based, low-data-rate sensors and 802.11-based, high-data rate devices. The Institute of Electrical and Electronics Engineers (IEEE) 802.11 protocol is used to form the wireless mesh between iNodes. The IEEE 802.15.4 (or 802.15.4 derivative protocols, such as ISA100.11a) communication protocol is used for point-to-point communication between a sensor node and neighboring iNode(s). With these protocols, a shared key is used for secure communication, but since a mesh network may involve multiple links in routing the communication, the same key is used to protect each link. There is a need to provide different keys for communication between any two

mesh network links, thus providing distinct link security through the routing network. A protocol for establishing distinct link keys is presented in Section 4.

For enhanced reliability, the network supports redundant communication paths between nodes. Data is injected into the network from leaf, device or client, nodes, such as sensors, and transported within the mesh network through the iNodes. This architecture provides versatile, multifunctional, plant-wide coverage for wireless sensing, asset location tracking, personnel tracking, and communications.

Industrial plants value wireless networks mainly for cost savings in the installation and maintenance of wires. Moreover, the reduced cost of wiring translates into more sensor and sensing points in the plant, which may lead to higher safety and productivity. Wireless sensors can also be installed on rotating equipment and for temporary online experiments, process optimization, and safety. To realize these advantages, wireless networks must be reliable and secure.

2.1 MAC Layer Communication Features

The OneWireless MAC layer consists of two different mechanisms. A time-division, multiplexed scheme is used for the communication of sensor data from the sensor/actuator to the iNode. The Time-Division Multiple Access (TDMA) scheme allows for efficient use of battery and bandwidth resources. A sensor's communication electronics can wake up in their assigned time slots, send a message, and go back to sleep as soon as the receipt of the message is acknowledged, thereby saving battery power.

The iNodes use a Carrier Sense Multiple Access (CSMA), as prescribed in the 802.11 specification, scheme to communicate amongst themselves. In CSMA, every iNode contends for access to the medium; by listening before talking, the iNode senses if the medium is free, and, if so, transmits a message. Because iNodes are line-powered, they can afford to listen continuously and transmit if the medium is free. The number of iNodes will be much less compared to the number of sensor/actuator nodes in a plant. Fewer numbers of CSMA iNodes contending for the medium are easier to manage and provide efficient and effective communication in terms of throughput, reliability and latency of the data. Moreover, the radio frequency (RF) space used by the infrastructure is different from the frequency space used by the sensor nodes. Therefore, there is no interference between those two types of nodes.

2.2 Mitigations of Threats

The overall objective of the current work is to develop enhancements to a wireless mesh network for improved cost-effective, secure and reliable communications in industrial control systems. Table 1 summarizes the current state-of-the-art in wireless mesh networking and enhancements of security and reliability for various threats against wireless networks for Process Control System (PCS) applications. Subsequent sections describe some of these enhancements in more detail.

Threats	Current State-of-the-Art	Enhancements
PCS wireless network is unavailable due to outage of a given link.	Network is unavailable until routing protocol re-converges.	Simultaneous multi-path routing with link-disjoint paths mitigates against link outages, such as RF.
Compromised node must be excluded from the PCS wireless network.	All nodes must be re-keyed with new network-wide key.	Unicast traffic from compromised node can be rapidly excluded, while network-wide broadcast / multicast keys are updated.
New node must be securely added to PCS wireless network.	Network-wide key must be communicated via secure out-of-band channel to new node.	Certificate-based techniques allow secure add/drop of new nodes.
Safety-critical traffic gets dropped during plant emergencies.	IEEE 802.11e provides quality of service (QoS) for access point (AP)-based networks but not multi-hop networks.	Separate QoS policies tailored to emergencies are activated during a crisis situation.
Man-in-the-Middle Attack on Integrity.	Malicious traffic can be injected along one path. Masquerading is possible if a single network key is used. With a single network key, masquerading is not mitigated. An adversary can inject traffic at will and send it through many different routes.	Link-disjoint, multi-path routing requires adversary to subvert multiple paths. If the adversary can modulate on the channel, he can insert data. The only question is where it will be stopped. One would like it stopped at the first hop. If not there, it must be stopped at the receiving node. Without strong source identification, neither is possible.
Man-in-the-Middle Attack on Confidentiality.	Adversary can eavesdrop on existing links.	If there is a single network key, the adversary can potentially eavesdrop on all communication. The mitigation is the transition to per-link keys.

Threats	Current State-of-the-Art	Enhancements
Attribution	Unsupported.	Without non-repudiation, one cannot correctly attribute bad behavior. Distinct per-link keys makes attribution possible in some cases. Link keys positively identify participants so that if one behaves badly, within certain bounds, it will be noticed. Anyone can behave badly outside of those bounds and it may not be attributable.
Emergency responders lack situation awareness during plant emergencies.	Unsupported.	1. Secure add/drop of new nodes. 2. QoS prioritization. 3. Bandwidth-efficient multicast of situation awareness data.

Table 1 – Network Communication Threats vs. Mitigations

2.3 Infrastructure Security Benefits

To foster wide adoption of wireless networks that can bring cost advantages to the industry and safety features to first responders, a system has to be robust, secure, usable during a crisis, and able to display data in a succinct and timely manner. The following security mechanisms and benefits can be provided to an industrial sensor and actuator network comprised of both wired and wireless devices.

- Enhanced security via node registration, authentication to the network, and per-link authenticated encryption
- Availability by utilizing robust, redundant, non-overlapping routes along with varying levels of Quality of Service. (QoS)
- Availability of SA data to first responders in the event of an emergency.

In addition to enhancing the commercially available OneWireless industrial wireless network, these features can be applied generically to other wireless mesh communication products. The result will be a system capable of reliable and secure wireless communication for control of industrial processes. The system will offer mechanisms for confidentiality, authenticity, and robustness to denial of service attacks, thus offering improved resilience against external cyber attacks. Simultaneously, the system will provide a secure and seamless on-the-fly integration of trusted entities, such as an emergency first responder communications network.

3 Network Key Management

The use of wireless mesh networks by critical infrastructure utilities implies the need for secure data transmissions. Communication systems that traditionally needed little security are becoming targets for a wider range of interested malicious parties. Further, many traditional communication systems are evolving into a more networked and sometimes wireless approach. This evolution is a double-edged sword. There are many functional benefits to having a large collection of devices networked together with high bandwidth wireless links, e.g, the ability to quickly communicate and share resources and information between mobile devices without a significant physical infrastructure. However, this larger degree of connectivity makes for a growing number of accessible targets that simply did not exist before with PCS applications.

To combat the vulnerabilities that accompany this evolution in the communications environment, cryptographic security measures can be applied to shore up the defenses. The measures tend to center around confidentiality, integrity, and availability, which is often called the "security CIA triad." The network design must allow the authorized communications to occur while denying access to unauthorized entities. The network should also be robust enough so that it will function even in the face of adversarial manipulations. Finally, it must function and be manageable even given the performance costs and penalties associated with the measures that provide confidentiality and integrity.

In general, two primary cryptographic methods are employed to provide confidentiality and integrity. They are distinguished by the type of keys used. Public key methods tend to be computationally intensive, but they have a number of very powerful benefits including no need to keep the public key secret, digital signatures, and non-repudiation. On the other hand, symmetric key methods are the workhorse of most secure communication systems. They are computationally efficient and their usage is the end product of many public key operations.

It is important to understand that ensuring, or even measuring, the security of a system is impossible to perform fully [1]. Therefore, the system designers and owners must use risk management. When designing a secure communication system, one must consider the threats against the system as well as the impact of the security mechanisms applied. Those impacts generally come in two flavors: operational and managerial. Security measures have an operational cost, which tend to be more easily quantified than managerial costs. Operational costs include processing cycles, bandwidth overhead, and delays associated with retrieval of information from a secure source. The costs associated with the management of the security measures tend to be harder to quantify and are often either overlooked or else drive the basic security profile. Depending on the system and the level of security that is required, more managerial costs are usually incurred. In an effort to show how some of these operational and managerial tradeoffs come to bear on network security, this document includes a discussion of three different security paradigms.

3.1 Single Network Key with Manual Rekey

This simple scenario is to provide each node in the network with a single key for confidentiality and integrity. Operationally, this approach requires each node to be able to encrypt and authenticate data. From a management perspective, it is the simplest method. The network key, which is usually a symmetric key, must be installed in each device before that device may participate in the secure communications. Once the key is installed, that device may communicate freely within the network.

Since this method uses a manual rekey, operationally there is very little overhead associated with the security. The traffic flows are simple and efficient, with no security overhead/maintenance messages involved in the system. The hardware and software needed to support the method is as minimal as can be expected. The device managing the installation of the keys can be extremely simple. It need only have a method to create and store a key and then transfer that key when needed. Since there must be a direct link between the two, the communication protocol between the key management device and the network device can be simple as well.

This approach to security, having a single network/group key, is simple to design and maintain and can be amazingly effective against a large number of the most common threats. Of course we assume that the system has been properly designed and implemented. If so, an adversary that does not know the network key must be able to break the cryptography that secures the data or find some other way of determining the key. If the cryptographic methods are strong, then cryptanalysis of the key is very difficult if not computationally infeasible.

However, as described, this simple method of security has a number of operational and security drawbacks in that it is rigid and inflexible. A compromise of a single node may have catastrophic consequences. Recovery from compromise requires each device in the field to be physically handled and reinitialized before the network may operate securely again.

When a new node comes into the network and is given the network key, that node now has the key for all prior traffic. In some cases, this may be acceptable. In others, it may be undesirable. When a node leaves the network, it likely has the network key still and will be able to recover any and all traffic. Even if the device has limited communication ability, it will still have the key. For an adversary without a key, picking through retired network communication boxes would be an excellent place to acquire the key.

With a single network key, no node has an individual identity. Even if each node is given a unique identifier, the key is what provides the security so a misbehaving node may masquerade as any other node. Therefore one can never be certain that bad behavior has been attributed correctly to a particular ill-behaving node. Because there is no unique, secure identity, an adversary may have free reign in the system, influencing it to behave in any way that the system is physically able to support. Finally, positive removal of any particular ill-behaving node is tantamount starting the network over from scratch.

3.2 Single Network Key with Electronic Rekey

Given some of the downsides of the method above, combined with the fact that cryptographic best practices limit the lifetime for which a key should be used, most networks require some method to change keys electronically rather than manually. A node in the network, called the key manager (KM) is usually designated to manage all keys in the network. From a protocol standpoint, securely bringing a new member into the network is easy if one has a way to distribute new keys to every other member. When a new member joins, a rekey message can be sent to all existing members before the new node is added. Then the new node is given the new key. A very simple method of doing the distribution is analogous to encrypting the new key with the old. This simple method of admitting new members is not without its flaws. For instance, any node may assume the role of initiation and admit a new member. If the KM is offline when this happens, the KM will afterward be out of sync with the rest of the network and will not be able to perform its duties. Then a manual rekey will be required as described in the previous subsection.

This problem can be addressed as follows. To remove a leaving node and prevent false rekeys, each node must have a unique rekey key that is known only by the KM and the node. Via this unique key, the KM may send a secure rekey message to the node with which the key is associated. This somewhat "out-of-band" message will inform existing nodes of a join or leave and provide them with a new network key to continue to protect the base traffic flows.

It should be noted that by allowing rekey messages within the network, this effectively separates the network into two distinct layers. The two layers use the same physical medium, but logically they are separate networks. One layer/network carries the normal network traffic and the other handles the rekey messages. So, even though a node has a key that allows it to be distinguished from other nodes by the KM, this key may not secure network traffic and so may not allow the node to be distinguished from other nodes when viewing things from a network traffic perspective. In particular, at any given time, since all network traffic is secured by a single key, the network still has the problem with the attribution of misbehaving nodes.

Providing a mechanism for electronic rekey solves a couple of security issues and mitigates the typically prohibitive requirement for manual rekey. However, electronic rekey does come with costs of its own. The network layering spoken of earlier may be viewed as a resource consumption issue. A certain percentage of the network resources will be consumed by the key management messages. Thus, the network capacity that may be devoted to the actual network traffic will be reduced by some percentage. If the frequency of network joins and leaves is high and the network is large, the percentage of consumed resources may be significant. Several protocols, [2] [3] among others, have been developed to help reduce the bandwidth consumption required for rekey in large networks.

The protocols and procedures required to cover all the rekey bases are not simple or easy to implement. Each network node is required to be able to seamlessly transition from network messages to key management messages and process each properly. The KM must be able to act as any network device. It must store the unique rekey keys and

be able to associate those with the correct node. It must also be powerful enough to organize and initiate the rekey process and follow it through to completion given a number of possible fault scenarios.

Given the hoops one must jump through to manage the network key and the rekey keys, the single network key is able to provide efficient cryptographic processing of the network data. However, the single network key still allows threats that stem from a lack of unique network identity.

3.3 Unique Network Identity and Multiple Network Keys

The change in complexity required to transition from a single network key with manual rekey to one with electronic rekey is surprisingly high. The baseline approaches given so far can be accomplished in their entirety with symmetric keys. Of course one may derive additional security and functional benefits by using public keys, but the baseline does not need them.

In order to give each node in the network a unique and cryptographically strong identity, public keys must be used. Being able to positively identify every node by its identifier as well as its key, makes it very difficult for rogue nodes to masquerade as another. This also means that correct attribution of misbehavior may be possible. With enforced link encryption, one may expect to limit the scope of a compromise. Given that a node, say Node A, is able to positively identify Node B, then with common public key exchange methods, Node A and Node B may create a link key that is unique to and known only by the pair. No other node may eavesdrop on the communications or inject false data between Node A and Node B as long as they are using the link key that is unique to them.

The inclusion of public keys into the design dramatically increases the management complexity. Similarly, the usage of public key methods dramatically impacts the operational overhead required for security. The typical public key operation is two to three orders of magnitude slower than that of the closest analogue in the symmetric key world. This is one reason that nodes create a symmetric link key. They use public key methods to establish a root of trust and then leveraging that root, they create a symmetric link key to allow efficient subsequent communications. Once the link keys are in place network communications are no more costly than if a single network key were to be used. The only minor difference is that the node must be able to store a link key for every link it maintains and then be able to use the right key at the right time to facilitate the communications. The use of per-link keys is problematic in the "mobile ad hoc networks" proposed for military applications. However, the infrastructure nodes used in PCS applications are typically stationary within a fixed, slowly varying neighbor set.

One of the common methods of providing identification for a node is through the use of digital certificates. The certificate is a message digitally signed by the KM, or another certificate authority recognized by the KM, to give permission for a node to operate in the network. Essentially, this certificate binds a node's identification with its public key. Once the permission is given, within the time constraints given within the certificate, all nodes within the network will accept the certificate and communicate with the node.

Since the KM must be able to remove nodes when needed, the KM must maintain a certificate revocation list (CRL). This list allows nodes to check to see if any other node has prematurely been removed from access to the network services. Verification of certificates and the checking of the CRL may be time consuming, and in a network with a large and dynamic membership, maintaining the CRL may be challenging. Making the CRL easily and efficiently accessible to network nodes is an open area of research. However, it should be manageable for PCS at a single facility since the membership should have infrequent changes.

4 MAC-Layer Security for Wireless Mesh Networks

Secure communication in the wireless mesh network is dependent upon the secrecy and strength of the keys used in cryptographic algorithms for confidentiality and integrity of data and for authentication of authorized users. Typical wireless-mesh networks use a single key at the Media Access Control (MAC) layer for the entire network. That key is deployed using the management console of the routing nodes. While very convenient, this approach has the security problems described in Sections 3.1 and 3.2. If any of the nodes in the mesh are compromised, then the network-wide key can be exposed and the entire network compromised. To avoid such a compromise, we utilize a public key certificate based approach that establishes separate, unique keys for each link in the wireless mesh. This provides a significant improvement in security over current approaches.

Each node in the wireless mesh will have a globally unique address to be used as an identifier for secure authentication purposes. A standard example is an address compliant with the Institute of Electrical and Electronics Engineers (IEEE) registration authority, commonly referred to as an IEEE address. The wireless network typically has a Central Trust Authority (CTA). Trust in network security is rooted in the CTA, which will perform Certificate Authority (CA) functions and is responsible for verifying the credentials of authorized nodes and for generating and maintaining certificates for the nodes.

The first step towards constructing the secure wireless network is the establishment of the CA to support a public key infrastructure. The CA must create a root certificate for itself and protect the secrecy of the associated private key. Knowledge of this private key by unauthorized parties will destroy the security of the entire system. The CA will then create digital certificates for each node that registers for access to the network with the proper credentials. The CA will establish and maintain a certificate database and list of revoked certificates. At deployment time, each new node will register its IEEE address with the CA and be issued a private/public key pair. A certificate with the node's public key and IEEE address is signed by the CA. The following is loaded into the node during registration.

- A copy of the CA's root certificate
- The node's certificate and its associated private encryption key
- A certificate revocation list (CRL)
- Synchronization of date and time

The new node is then deployed to cover the appropriate area within the facility. As it is brought online, it goes through its neighbor discovery process. After the discovery process, the new node will have a neighbor set. For each node in that set, the new node will exchange certificates with the neighbor. The new node and its neighbor will verify each other's identities based on the CTA's root certificate, asserting that the node has been seen and authorized by the CTA. Those nodes can then proceed to generate and exchange a symmetric key unique for that link. The process above is repeated for all the nodes in the neighbor set of the new node.

At the conclusion of the protocol, the new node has a set of independent secured links to each of its neighbors. Retrieving the appropriate cryptographic material for an incoming packet from a neighbor is a simple lookup. The same holds for outgoing packets to the next neighbor in the route. The lookup can be based on addresses since they are globally unique. The wireless network topology in an industrial system is relatively fixed. Thus, there is little overhead for link maintenance. For broadcast and multicast packets, a network or group wide key is used. The key is distributed by the CTA over a dedicated link between each node and the CTA. Upon removal of an node in the system, the CTA can redeploy the appropriate keys.

4.1 Key Management Protocol Outline for MAC-Layer Security

This section provides specifications of the key management protocol designed to satisfy the needs of MAC-layer security for distinct link authenticated encryption. Upon receiving a new node at a particular critical infrastructure site, it will be commissioned for use at that site by registering with the CA and receiving its own certificate and private decryption key. Access to the CA is assumed to be limited only to authorized personnel. Trust in the registration process is based on the personnel allowed to register nodes, the security of the CA computer system, and a direct, local, physical connection (i.e., a crossover cable) with the CA during registration. Once a node is installed in the mesh network, it is now ready to establish a MAC-layer encryption key with its neighbors. A separate, unique key will be used to communicate with each neighbor.

One node will initiate the key exchange protocol with another node. They will exchange and verify each other's certificates and use the public keys in the certificates to encrypt a random number that each node generates individually. They exchange the encrypted random numbers, decrypt them and hash the two random numbers together to create the shared encryption key. One advantage of this proposed protocol is that it does not require digital signature generation by either node, only signature verification of the certificates signed by the CA. Therefore nodes do not need their own RSA digital signature key pair. They only need the CA's public key for verification of certificates. The following protocol description is appropriate for a generic set of network nodes, including nodes. This approach is proposed since it reduces the computational requirements at the nodes.

Standards influencing the protocol include the following:

- Draft FIPS 186-3: Digital Signature Standard [4]
    - Refer to this standard for the RSA Digital Signature Algorithm
- PKCS #1 v2.1: RSA Cryptography Standard [5]
    - Refer to this standard for RSA Encryption/Decryption
- FIPS 180-2: Secure Hash Standard [6]
    - Refer to this standard for SHA-256
- RFC5280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile [7]
    - Refer to this standard for X.509 certificate and CRL operations
- FIPS 197: Advanced Encryption Standard [8]
    - Refer to this standard for AES Encryption/Decryption

All digital signature operations shall be performed with RSA per [4] using SHA-256 [6] as the hash algorithm. All RSA encryption/decryption operations shall be performed per [5]. All AES encryption/decryption operations shall be performed per [8]. All certificate and CRL operations shall be performed per [7].

Notation, Labels, Definitions:

CA – Certificate Authority, administered by the CTA

L(x) is the length of x in bits

$X_{S,CA}$, $Y_{S,CA}$ – CA's RSA Digital Signature Private key, Public key pair

- M = L(Signature Modulus) ≥ 1024
- 16 ✪ L($X_{S,CA}$) ✪ M
- L($Y_{S,CA}$) = M

$X_{E,i}$, $Y_{E,i}$ – Node i's RSA Encryption Private key, Public key pair

- L(Encryption Modulus) =M
- 16 ✪ L($X_{E,i}$,) ✪ M
- L($Y_{E,i}$) = M

$X_{S,CA}$ and $X_{E,i}$, – must be kept secret

K_{ij} – 128-bit AES Session Key shared by nodes i and j

RSA_E[k, z] – RSA Encryption of z using key k

RSA_D[k, z] – RSA Decryption of z using key k

$AES_E[k, z]$ – AES-128 Encryption of z using key k

$AES_D[k, z]$ – AES-128 Decryption of z using key k

SHA-256[m] – The 256-bit Secure Hash Algorithm hash of message m

SIG_i – RSA Digital Signature of Node i's certificate information, signed by the CA

- Length of SIG_i = M

CSN – Certificate serial number

$CERT_i$ – Node i's X.509 Version 3 certificate, identified by a CSN

D_{CERT} – Certificate Database indexed with CSNs

CRL – X.509 Certificate Revocation List of revoked CSNs

R – 128-bit random number

MSB-128[H] – The most significant 128 bits of H

Name	Value (Decimal)
CREDENTIAL_SUBMIT	7
CREDENTIAL_INIT	25
KEY_XCHG_START	30
KEY_XCHG_REPLY	42
KEY_XCHG_FINAL	45
KEY_XCHG_VERIFY	51

Table 2 – Message Types (1 Byte)

The following subsections provide design specifications of key management activities required to support the establishment of a unique symmetric key between Nodes i and j. Protocols 4.1.2, 4.1.3, and 4.1.4 are summarized in a data flow diagram shown in Figure 2. For each protocol, particularly 4.1.4, implementations should adhere to a finite state machine with strict enforcement of state transitions to ensure proper flow of the protocol.

4.1.1 Protocol: CA Establishment

1. CA creates its own RSA digital signature parameters, including its private/public RSA key pair, $[X_{S,CA}, Y_{S,CA}]$.
2. CA creates its own self-signed certificate, $CERT_{CA}$.
3. CA establishes a Certificate Database, D_{CERT}, based on CSN's.

4. CA establishes a CRL of all revoked node's certificates.

4.1.2 Protocol: Node Registration

1. Node i connects to the Certificate Registration port of the CA.
2. Node i synchronizes its time and date with the CA. The time and date are used by nodes to check that certificates are in the proper time window.
3. Node i sends to the CA: [CREDENTIAL_SUBMIT, Node i's Access Point MAC address, Node i's Mesh Radio MAC address].
4. CA generates RSA encryption/decryption parameters for Node i, including an encryption private/public key pair, [$X_{E,i}$, $Y_{E,i}$].
5. CA creates an X.509v3 certificate, $CERT_i$, for Node i.
6. CA stores $CERT_i$ in the Certificate Database.
7. CA sends to Node i: [CREDENTIAL_INIT, $X_{E,i}$, $CERT_{CA}$, $CERT_i$, CRL].

4.1.3 Protocol: Node Removal / Certificate Revocation

1. CA adds removed/revoked node's CSN information to the CRL.
2. CA digitally signs the new CRL.
3. CA broadcasts new CRL to all remaining nodes.

4.1.4 Protocol: Symmetric Key Exchange

1. Node i sends to Node j: [KEY_XCHG_START, $CERT_i$].
2. Node j checks to see if $CERT_i$ is in the proper time window.
3. Node j checks the CRL for $CERT_i$'s CSN.
4. Node j verifies $CERT_i$'s signature.
5. Node j generates a 128-bit random key, R_j.
6. Node j encrypts R_j with Node i's public key, $E_{Rj} = RSA_E[Y_i, R_j]$.
7. Node j sends to Node i: [KEY_XCHG_REPLY, $CERT_j$, E_{Rj}].
8. Node i checks to see if $CERT_j$ is in the proper time window.
9. Node i checks CRL for $CERT_j$'s CSN.
10. Node i verifies $CERT_j$'s signature.
11. Node i decrypts $R_j = RSA_D[X_i, E_{Rj}]$.
12. Node i generates a 128-bit random number, R_i.
13. Node i computes the shared key, K_{ij} = MSB-128[SHA-256[R_i, R_j]].
14. Node i encrypts R_i with Node j's public key, $E_{Ri} = RSA_E[Y_j, [R_i, AES_E[K_{ij}, R_i]]]$.

15. Node i sends to Node j: [KEY_XCHG_FINAL, E_{Ri}].
16. Node j decrypts $R_i = RSA_D[X_j, E_{Ri}]$.
17. Node j computes the shared key, K_{ij} = MSB-128[SHA-256[R_i, R_j]].
18. Node j verifies R_i by comparing it with $AES_D[K_{ij}, R_i]$.
19. Node j hashes K_{ij}, H_K = MSB-128[SHA-256[K_{ij}]].
20. Node j encrypts H_K with AES-128 and K_{ij}, $E_{HK} = AES_E[K_{ij}, H_K]$.
21. Node j sends to Node i: [KEY_XCHG_VERIFY, E_{HK}].
22. Node i decrypts $H_K = AES_D[K_{ij}, E_{HK}]$.
23. Node i verifies that H_K = MSB-128[SHA-256[K_{ij}]].

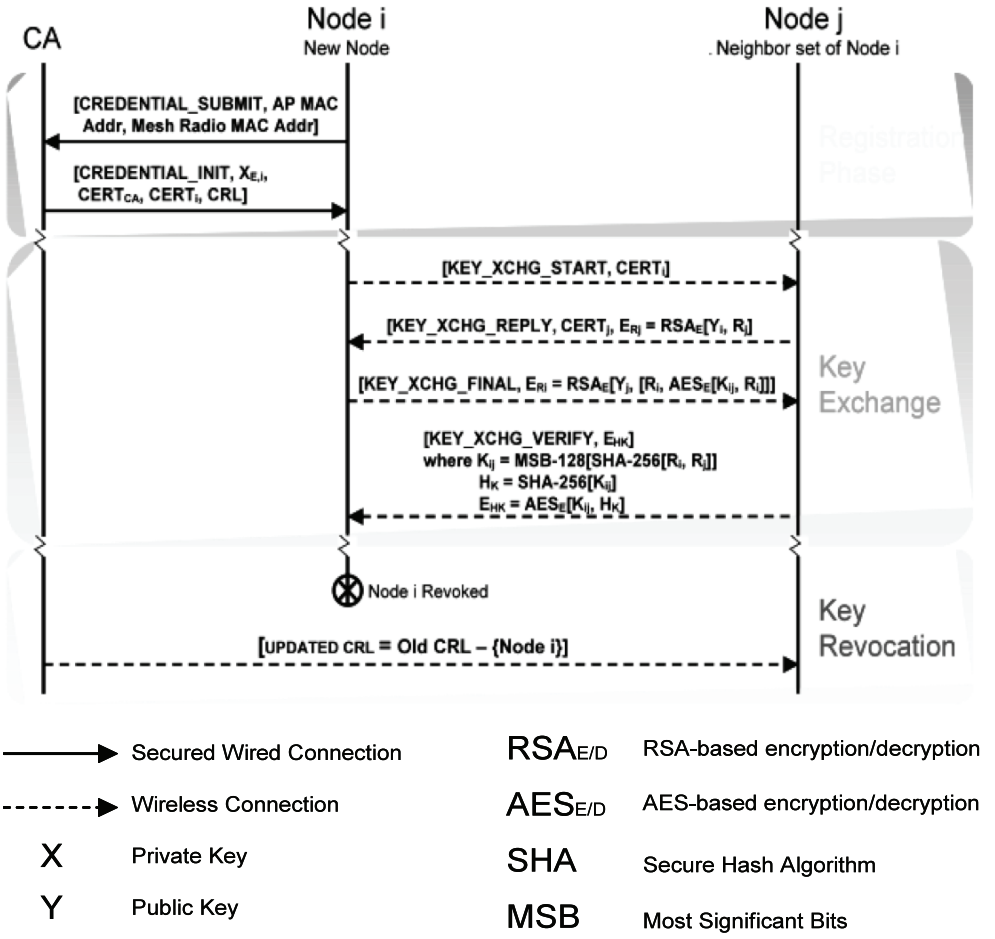

Figure 2 – Data Flow Diagram for Node Registration, Key Exchange & Revocation

5 First Responder Credentialing

During an emergency, first responders (FR's) need access to the wireless mesh network for the purpose of maintaining situation awareness. However, FR's should not have access to the network during non-emergency periods. This section presents a design to allow FR's access to the mesh network with their own or the facility's equipment when they arrive on-site in response to an emergency and to terminate that access when the emergency is over or sooner if necessary. Recall that the Honeywell OneWireless network is composed of a mesh of nodes, each possessing a radio that is used to allow FR's with 802.11b/g devices (such as laptops or PDAs) to connect to the mesh network. The protocols presented below address only the problem of time-limiting the access of FR's to the wireless mesh network and assume that communication security is handled via the Remote Authentication Dial In User Service (RADIUS) and 802.1x protocols.

In an emergency, having access to the network in the same manner as the wireless workers will give first responders an improved situation awareness. For example, the emergency responders can learn the location and status of key assets as well as hazardous materials. However, distributing new credentials in a high-stress environment can be error prone. We propose to deploy a system where the credentials for FR authentication to the RADIUS Server (RS) would be preloaded on a digital media connected to the CTA. Those credentials would be time-limited, but the clock would not start until the media was removed from the trust center.

When first responders need access to the network, the software on their devices can use the provided credentials to authenticate the device to the RS. They will then gain access to the network resources following a predefined access policy. The QoS properties of the network ensure priority to FR's in this situation. The traffic belonging to FR's may be differentiated through the use of a virtual LAN, such as a separate SSID, within the network. Once the FR's have completed their mission, their credentials will be revoked or their privileges rendered useless to avoid a security leak.

There is a tradeoff between usability and security in not overburdening the FR's with security responsibilities during the emergency. A Storage Device (SD), such as a memory card or a USB drive, holding credentials for use in authenticating to the network will be issued to all registered FR's. The SD can be installed in a FR's own equipment or in equipment issued by the facility. The only typing required of FR's is a login password to the computing device, either their own or the facility's, which can be eliminated if the facility so desires. The FR Credentialing Protocol consists of a set of Preparation steps and a set of Emergency Operation steps as follows.

5.1 Protocol: Preparation

1. The Plant Manager (PM) at the host facility must establish and maintain a FR Database (FRD) containing a list of approved FR's that register with the facility by presenting proper credentials[2].

[2] Credentials required for proper identification to be decided upon by the facility.

2. The RS must have access to the FRD and will associate access privileges with each FR in the database.
3. FR's that respond to an emergency, but are not in the database, must first register with the PM before they can receive credentials and access the network.
4. The PM shall acquire SD's to be issued to each registered FR upon their arrival at the facility during an emergency. The SD's shall be loaded with a username and password as soon as possible after an emergency begins or in advance for those FR's that are preregistered.
5. The PM (or delegate) at the host facility will manage the privileges of the FR's. The following privileges are specified regarding access to the wireless mesh network.
6. No Access – DEFAULT setting.
7. Full Access.

5.2 Protocol: Emergency Operation

1. As soon as possible after the beginning of an emergency,
   (a) Each FR SD shall be loaded with a unique username and a unique strong password.
   (b) The permissions for all approved FR's shall be set to Full Access.
2. As soon as possible after the beginning of an emergency, the RS shall set a count-down timer, T, to 72 hours[3] and begin counting down. The PM can extend the duration of the FR privileges at any time (the count-down timer shall increase by the amount of time extended).
3. Once on site in response to an emergency, the FR will be given a SD to install in his own computing device or be given a computing device with a SD already installed. The following information must be added to the database for each FR issued a SD.
   (a) FR credential
   (b) Computing device identifier (e.g., Wireless network card's MAC address)
   (c) SD identifier
   (d) Username
   (e) Password
4. Immediately after issuing to an FR, all equipment should be tested for access to the wireless mesh network.

[3] Or whatever is reasonable, based on the type of facility.

5. As an example, all FR's given a computing device during an emergency will be issued a random, unique, 3-digit PIN with which to login to the computing device. The PM must establish a user account on the computing device requiring the PIN for login access.

   (a) If the burden of a PIN is determined to be great, the PM can eliminate the PIN (i.e., not require a login password to the computing device) or assign all computing devices the default PIN of 000.

   (b) Longer PINs or stronger types of "two factor authentication" could be implemented if higher security is needed.

6. For each computing device offered by the facility to FR's, the PM must have an administrator account on the device.

7. Each node in the wireless mesh network will operate as a wireless access point and Network Access Server (NAS) and interact with FR clients.

8. When a FR attempts to connect to the network, the username and password stored on its SD shall be communicated automatically to the wireless access point with the highest signal strength. The NAS will submit an Access-Request message to the RS.

9. The RS will check the credentials in each connection request against information in the FRD and grant or deny the request based on current privilege settings and the following conditions.

   (a) The RS sends an Access-Reject response to the requesting NAS if any of the following conditions hold.

      - The count-down timer is zero.
      - All conditions of the Access-Request are not met.
      - The privilege setting for the FR is No Access.

   (b) The RS sends an Access-Accept response to the requesting NAS if all of the following conditions hold.

      - The count-down timer is greater than zero.
      - All conditions of the Access-Request are met.
      - The privilege setting for the FR is Full Access.

10. When T = 0, the RS will initiate termination of each FR session using a Packet of Disconnect (see RFC 3576, "Dynamic Authorization Extensions to RADIUS"). The RS will send to each NAS that is currently acting as the access point for a FR the Packet of Disconnect message. The NAS will terminate its session with that FR.

11. The PM can set FR privileges to No Access at any time. The PM shall set all FR privileges to No Access at the end of the emergency, regardless of how much

time has elapsed. The PM can initiate termination of all FR sessions at any time via the RS using a Packet of Disconnect (see RFC 3576, "Dynamic Authorization Extensions to RADIUS"). The RS will send to each NAS that is currently acting as the access point for a FR the Packet of Disconnect message. The NAS will terminate its session with that FR.

5.3 Benefits

By providing an intuitive mechanism for credential dissemination, we can reduce the time required to distribute cryptographic information during normal plant operations. At the same time, we provide a fast, effective way for first responders to get access to network resources from the wireless mesh of a critical infrastructure. Once the network is accessible, the first responders can retrieve the infrastructure's data and filter the portion that can be critical to their mission and personal safety, such as detection of toxic chemicals in a given area. The system also enables an intuitive method to revoke distributed credentials by reinserting in the CTA's interface to "clear" the keys. A built-in safety mechanism automatically revokes distributed credentials, thus preventing sensitive information from leaking out of the network after the emergency is over. Additional levels of protection can be implemented in the odnes or PCS to limit the data to which FR's have access. In addition, the nodes or PCS may limit what commands the FR's can issue. This paper's Concept of Operations (CONOPS) envisions that the FR's can view select PCS data during a plant emergency but not change any PCS settings.

6 Conclusions

Wireless communication networks are increasingly popular in industrial applications and critical infrastructure for functional and financial reasons. Effective security, however, continues to be an issue, given the bandwidth and computational limitations in some PCS environments. As such, this paper presents a computationally-feasible public key-based key exchange protocol that supports MAC-layer link authenticated encryption. Use of distinct per-link keys for communication between each pair of nodes in a network avoids the catastrophic problem of key compromise when a single network key is used. Key storage at each node is limited to a fraction of the total number of keys of neighbors only need to be stored.

A protocol for enforcing time-limits on first responder credentials for emergency access to a plant network is also presented. First responders will gain access to the wireless mesh network upon arrival at the emergency and registration with the plant's CTA. A RADIUS server will terminate first responder network sessions when the emergency is over or the authorized time has elapsed. The amount of time a first responder is allowed on the network can be extended easily without interaction. These security enhancements will improve the reliability and security for process monitoring and control of existing wireless mesh network systems used in critical infrastructure facilities.

About the Authors – Tim Draelos, Mark Torgerson, and Pete Sholander are members of the technical staff at Sandia National Laboratories.

Tim Draelos conducts research and development in cryptographic applications, implementations, and algorithms and has led and participated in several R&D projects in SCADA security. He holds a Ph.D. in electrical engineering from the University of New Mexico and is a member of the IEEE.

Mark Torgerson holds a Ph.D. in mathematics from the University of Arizona and is involved in research and development of cryptographic algorithms, analysis, and implementations. Most recently, he participated in the development of the SANDstorm algorithm for the NIST hash competition.

Pete Sholander holds a Ph.D. in Electrical Engineering from the Georgia Institute of Technology and has over 25 years of experience in communication and network security research and development.

Soumitri Kolavennu is senior research scientist at Honeywell International, as was Denis Foo Kune until he recently began a Ph.D. program at the University of Minnesota. Soumitri holds a Ph.D. in chemical engineering from Florida State University. Both Soumitri and Denis have design and implementation experience with wireless mesh network communication security, including with the OneWireless Industrial Wireless Sensor Network.

References

[1] M. Torgerson, "Security Metrics for Communication Systems," 12th International Command and Control Technology Symposium, June, 2007.

[2] D. Wallner, E. Harder, R. Agee, "Key Management for Multicast: Issues and Architectures," Request for Comments: 2627, June 1999.

[3] J. Opper, B. DeCleene, M. Leung, "Gateway Subset Difference Revocation," 2nd International Workshop on Wireless and Sensor Network Security, October 2006.

[4] National Institute of Standards and Technology (NIST), "Digital Signature Standard (DSS). Draft FIPS Publication 186-3," March, 2006.

[5] RSA Laboratories, "PKCS #1: RSA Cryptography Standard," version 2.1. http://www.rsa.com/rsalabs/node.asp?id=2125.

[6] National Institute of Standards and Technology (NIST), "Secure Hash Standard (SHS), FIPS Publication 180-2," January 27, 2000.

[7] RFC5280, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile."

[8] National Institute of Standards and Technology (NIST), "Advanced Encryption Standard (AES), FIPS Publication 197," November 26, 2001.