NENA Next Generation 9-1-1 Security (NG-SEC) Information Document

NENA NG9-1-1 Security Information Document

NENA-INF-015.1-2016
DSC Approval: 08/31/2016
PRC Approval: 11/18/2016
NENA Executive Board Approval: 12/08/2016
Next Scheduled Review Date: 12/08/2017

Prepared by:
National Emergency Number Association (NENA) Interconnection and Security Committee, Security Topics Subcommittee, NG-SEC Working Group

Published by NENA
Printed in USA

NENA
INFORMATION DOCUMENT
NOTICE

This Information Document (INF) is published by the National Emergency Number Association (NENA) as an information source for the designers, manufacturers, administrators and operators of systems to be utilized for the purpose of processing emergency calls. It is not intended to provide complete design or operation specifications or parameters or to assure the quality of performance for systems that process such equipment or services.

NENA reserves the right to revise this Information Document for any reason including, but not limited to:

- Conformity with criteria or standards promulgated by various agencies,
- Utilization of advances in the state of the technical arts,
- Or to reflect changes in the design of equipment, network interfaces or services described herein.

This document is an information source for the voluntary use of communication centers. It is not intended to be a complete operational directive.

It is possible that certain advances in technology or changes in governmental regulations will precede these revisions. All NENA documents are subject to change as technology or other influencing factors change. Therefore, this NENA document should not be the only source of information used. NENA recommends that readers contact their 9-1-1 System Service Provider (9-1-1 SSP) representative to ensure compatibility with the 9-1-1 network, and their legal counsel to ensure compliance with current regulations.

Patents may cover the specifications, techniques, or network interface/system characteristics disclosed herein. No license expressed or implied is hereby granted. This document shall not be construed as a suggestion to any manufacturer to modify or change any of its products, nor does this document represent any commitment by NENA or any affiliate thereof to purchase any product whether or not it provides the described characteristics.

This document has been prepared solely for the use of 9-1-1 System Service Providers, network interface and system vendors, participating telephone companies, 9-1-1 Authorities, etc.

By using this document, the user agrees that NENA will have no liability for any consequential, incidental, special, or punitive damages arising from use of the document.

NENA's Committees have developed this document. Recommendations for change to this document may be submitted to:

National Emergency Number Association
1700 Diagonal Rd, Suite 500
Alexandria, VA 22314
202.466.4911
or commleadership@nena.org

ACKNOWLEDGEMENTS

The National Emergency Number Association (NENA) Interconnection and Security Committee, Security Topics Subcommittee, NG-SEC Working Group developed this document.

NENA recognizes the following industry experts and their employers for their contributions in development of this document.

Executive Board Approval Date 12/08/2016

Members	**Employer**
Nate Wilcox, Interconnection and Security Committee Co-Chair	Emergicom, LLC
Steve O'Conor, ENP, Interconnection and Security Committee Co-Chair	Synergem Technologies, Inc.
Brian Knueppel, Security Topics Subcommittee Chair	Oracle (Acme Packet)
Patrick Voigt, ENP, Security Topics Subcommittee, Working Group Leader	Synergem Technologies, Inc.
Brian Rosen	Neustar
John Skain	Clinton County IL
Mike Vislocky	Network Orange
Steve Lagreid	King County WA
Roger Marshall	Comtech TCS
Christian Militeau, ENP	West Safety Services
Carl Rodabaugh, ENP	Midland County MI

Special Acknowledgements:

Delaine Arnold, ENP, Committee Resource Manager, has facilitated the production of this document through the prescribed approval process.

The NG-SEC Working Group is part of the NENA Development Group that is led by:

- Pete Eggimann, ENP, and Jim Shepard, ENP, Development Steering Council Co-Chairs
- Roger Hixson, ENP, Technical Issues Director
- Chris Carver, ENP, PSAP Operations Director

Table of Contents

1 Executive Overview

This information document is a companion to NENA-STA-010, Detailed Functional and Interface Specification for the NENA i3 Solution. Users should also refer to NENA 75-001 and 75-502. To effectively use this document, the user should have a clear understanding of the concepts and procedures described therein. This document provides detail of the mechanisms and best practices relative to security of the i3 system. This document describes procedures and best practices on how to deploy security for the system.

The 9-1-1 system contains a significant amount of sensitive data and communications. It is a likely target of deliberate attack. Security is not meant to be convenient, rather it is meant to protect.

Unfortunately, somewhere in our country, NG9-1-1 will be attacked. It is not a matter of if but when, where and for how long. This has been informally echoed by the NENA NG-SEC (Security) Working Group over the years during work sessions that are comprised of subject matter experts. Network security should be considered the most critical component of NG9-1-1 by everyone – PSAPs and vendors alike. Network and system security should be considered ahead of any new product solution or technology. The mindset and posture of all stakeholders in NG9-1-1 must change from "closed network" to a constant "threat assessment". No solution can be bullet-proof but we must make every effort to defend in multiple layers, follow industry best practices, NENA standards and guidance and collaborate together as a community.

2 Introduction

2.1 Operations Impacts Summary

This NENA Information Document may have impacts on ESInet, NGCS and PSAP architecture and construct.

2.2 Technical Impacts Summary

This NENA Information Document will have impacts on technical aspects of the NG9-1-1 industry, particularly with the 9-1-1 network (including data) and/or CPE equipment, and most specifically with i3 Functional Elements.

2.3 Security Impacts Summary

This is a security information document which may impact other NENA standards and should be reviewed by each committee.

2.4 Reason for Issue/Reissue

NENA reserves the right to modify this document. Upon revision, the reason(s) will be provided in the table below.

Doc #	Approval Date	Reason For Changes
NENA-INF-015.1-2016	12/08/2016	Initial Document

NENA-INF-015.1-2016	05/07/2024	Archived

2.5 Recommendation for Additional Development Work

This information document does not require additional standards or development. Instead, it makes suggestions relative to existing standards.

2.6 Anticipated Timeline

Applicable sections of this information document should be implemented at initial deployment.

2.7 Cost Factors

This information document refers to standards which will have a cost impact to the entities that deploy security mechanisms. Security needs to be part of the Authority's NG9-1-1 initial and ongoing budget. The cost of implementing security is high. The cost of dealing with a security incident is much higher.

2.8 Cost Recovery Considerations

Normal business practices shall be assumed to be the cost recovery mechanism.

2.9 Additional Impacts (non-cost related)

The information or requirements contained in this NENA document are expected to be relative to other NENA documentation, based on the analysis of the authoring group. The primary impacts are expected to include:

- Process
- Policy
- General security
- Encryption

2.10 Intellectual Property Rights (IPR) Policy

NOTE – The user's attention is called to the possibility that compliance with this document may require use of an invention covered by patent rights. By publication of this document, NENA takes no position with respect to the validity of any such claim(s) or of any patent rights in connection therewith. If a patent holder has filed a statement of willingness to grant a license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license, then details may be obtained from NENA by contacting the Committee Resource Manager identified on NENA's website at www.nena.org/ipr.

Consistent with the NENA IPR Policy, available at www.nena.org/ipr, NENA invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this document.

Please address the information to:

National Emergency Number Association
1700 Diagonal Rd, Suite 500
Alexandria, VA 22314
202.466.4911
or commleadership@nena.org

2.11 Abbreviations, Terms and Definitions

See NENA-ADM-000, NENA Master Glossary of 9-1-1 Terminology, located on the NENA web site for a complete listing of terms used in NENA documents. All abbreviations used in this document are listed below, along with any new or updated terms and definitions.

Term or Abbreviation (Expansion)	Definition / Description
ACL (Access Control List)	A configurable permit/deny list at Layers 3 and 5.
B2BUA (Back-to-Back User Agent	A back-to-back user agent is a SIP element that relays signaling mechanisms while performing some alteration or modification of the messages that would otherwise not be permitted by a proxy server. A logical entity that receives a request and processes it as a user agent server (UAS). In order to determine how the request should be answered, it acts as a user agent client (UAC) and generates requests. Unlike a proxy server it maintains dialog state and must participate in all requests sent on the dialogs it established.
Blackholing	IP address where incoming or outgoing data packets are silently discarded without informing the source that the data did not reach its intended route
BCF (Border Control Function)	Provides a secure entry into the ESInet for emergency calls presented to the network. The BCF incorporates firewall, admission control, and may include anchoring of session and media as well as other security mechanisms to prevent deliberate or malicious attacks on PSAPs or other entities connected to the ESInet.
CA (Certificate Authority)	A trusted entity that issues electronic documents that verify a digital entity's identity on the Internet. The electronic documents, which are called digital certificates, are an essential part of secure communication and play an important part in the public key infrastructure (PKI).

Term or Abbreviation (Expansion)	Definition / Description
CERT (Computer Emergency Readiness Team)	Information technology (IT) security organization. The purpose of CERT is to respond to computer security incidents, report on vulnerabilities and promote effective IT security practices throughout the country. https://www.us-cert.gov/
CP/CPS (Certificate Policy/Certification Practice Statement)	Practice and Policy over the PKI. A public key infrastructure (PKI) supports the distribution and identification of public encryption keys, enabling users and computers to both securely exchange data over networks such as the Internet and verify the identity of the other party.
CRL (Certificate Revocation List)	Certificate Revocation List (CRL) is one of two common methods when using a public key infrastructure for maintaining access to servers in a network. The other, newer method, which has superseded CRL in some cases, is Online Certificate Status Protocol (OCSP).
CSP (Cryptographic Service Provider)	A library that provides cryptographic functions
DoS (Denial of Service)	An incident in which a user or organization is deprived of the services of a resource they would normally expect to have. In a distributed denial-of-service, large numbers of compromised systems (sometimes called a botnet) attack a single target.
DMZ (Demilitarized Zone)	In computer networks, a DMZ (demilitarized zone) is a physical or logical sub-network that separates an internal local area network (LAN) from other untrusted networks, usually the Internet.
DDoS (Distributed Denial of Service)	A denial of service (DoS) attack is an incident in which a user or organization is deprived of the services of a resource they would normally expect to have. In a distributed denial-of-service, large numbers of compromised systems (sometimes called a botnet) attack a single target.
DNS (Domain Name System)	A globally distributed database for the resolution of host names to IP addresses.

Term or Abbreviation (Expansion)	Definition / Description
DNSSEC (Domain Name System Security Extensions)	DNS Security Extensions (DNSSEC) are a set of Internet Engineering Task Force (IETF) standards created to address vulnerabilities in the Domain Name System (DNS) and protect it from online threats.
DRM (Data Rights Management)	Wherever data that might be considered sensitive, which in 9-1-1 is nearly all data, should be subject to data rights management. This involves careful consideration of agent roles, and construction of the DRM rule sets, secure credential handling and appropriate handling of errors and alarms.
EIDD (Emergency Incident Data Document)	A National Information Exchange Model (NIEM) conformant object that is used to share emergency incident information between and among authorized entities and systems.
ESInet (Emergency Services IP Network)	An ESInet is a managed IP network that is used for emergency services communications, and which can be shared by all public safety agencies. It provides the IP transport infrastructure upon which independent application platforms and core services can be deployed, including, but not restricted to, those necessary for providing NG9-1-1 services. ESInets may be constructed from a mix of dedicated and shared facilities. ESInets may be interconnected at local, regional, state, federal, national and international levels to form an IP-based inter-network (network of networks). The term ESInet designates the network, not the services that ride on the network. See NG9-1-1 Core Services.

Term or Abbreviation (Expansion)	Definition / Description
FIPS (Federal Information Processing Standards)	Federal Information Processing Standards (FIPS) is a standard for adoption and use by United States Federal departments and agencies that has been developed within the Information Technology Laboratory and published by the National Institute of Standards and Technology (NIST), a part of the U.S. Department of Commerce. FIPS describe document processing, encryption algorithms and other information technology standards for use within non-military government agencies and by government contractors and vendors who work with the agencies. The standards cover a specific topic in information technology (IT) and strive to achieve a common level of quality or interoperability.
FQDN (Fully Qualified Domain Name)	The complete domain name for a specific computer, or host, on the Internet.
IPSec (Internet Protocol Security)	IPsec (Internet Protocol Security) is a framework for a set of protocols for security at the network or packet processing layer of network communication. IPsec provides two choices of security service: Authentication Header (AH), which essentially allows authentication of the sender of data, and Encapsulating Security Payload (ESP), which supports both authentication of the sender and encryption of data as well.
IPv4 (Internet Protocol version 4)	The fourth version of the Internet Protocol; uses 32-bit addresses.
IPv6 (Internet Protocol version 6)	The most recent version of the Internet Protocol; uses 128-bit addresses.
LNG (Legacy Network Gateway)	An NG9-1-1 Functional Element that provides an interface between a non-IP originating network and a Next Generation Core Services (NGCS) enabled network.

Term or Abbreviation (Expansion)	Definition / Description
LSRG (Legacy Selective Router Gateway)	The LSRG provides an interface between a 9-1-1 Selective Router (including legacy ALI), and an NGCS, enabling calls to be routed and/or transferred between Legacy and NG networks. A tool for the transition process from Legacy 9-1-1 to NG9-1-1.
MITM (Man in the Middle)	An attack method that allows an intruder to access sensitive information by intercepting and altering communications between the user of a public network and any requested element.
NAPT (Network Address Port Translator)	A methodology of remapping one IP address and port into another by modifying network address information in Internet Protocol (IP) datagram packet headers while they are in transit across a traffic routing device.
NAS (Network Attached Storage)	A type of dedicated file storage device that provides local-area network local area network (LAN) nodes with file-based shared storage through a standard Ethernet connection.
NAT (Network Address Translation)	NAT maps a single public address to one, or many internal addresses and all network IP addresses on the connected computers are local and cannot be seen by the outside world.
NGCS (NG9-1-1 Core Services)	The base set of services needed to process a 9-1-1 call on an ESInet. Includes the ESRP, ECRF, LVF, BCF, Bridge, Policy Store, Logging Services and typical IP services such as DNS and DHCP. The term NG9-1-1 Core Services includes the services and not the network on which they operate. See Emergency Services IP Network.
OCSP (Online Certificate Status Protocol)	OCSP (Online Certificate Status Protocol) is one of two common schemes for maintaining the security of a server and other network resources. The other, older method, which OCSP has superseded in some scenarios, is known as Certificate Revocation List (CRL).
PAT (Port Address Translator)	PAT or NAPT is an extension to NAT in that PAT uses TCP/UDP ports in addition to network addresses (IP addresses) to map many private network addresses to a single outside address.

Term or Abbreviation (Expansion)	Definition / Description
PCA (PSAP Credentialing Agency)	The root authority designated to issue and revoke security credentials (in the form of an X.509 certificate) to authorized 9-1-1 agencies in an i3-compliant infrastructure.
PKI (Public Key Infrastructure)	A set of hardware, software, people, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates and manage public-key encryption.
QoS (Quality of Service)	As related to data transmission, a measurement of latency, packet loss and jitter.
RTCP (Real-time Transport Control Protocol)	RTCP is a sister protocol of RTP and provides out-of-band control information for an RTP flow. It partners RTP in the delivery and packaging of multimedia data, but does not transport any data itself. It is used periodically to transmit control packets to participants in a streaming multimedia session. The primary function of RTCP is to provide feedback on the quality of service being provided by RTP. It gathers statistics on a media connection and information such as bytes sent, packets sent, lost packets, jitter, feedback and round trip delay. An application may use this information to increase the quality of service perhaps by limiting flow, or maybe using a low compression codec instead of a high compression codec. RTCP is used for Quality of Service (QoS) reporting.
RTP (Real-time Protocol)	An IP protocol used to transport media (voice, video, text) which has a real time constraint.
SAML (Security Assertion Markup Language)	An XML-based, open-standard data format for exchanging authentication and authorization data between an identity provider and another party.
SBC (Session Border Controller)	A commonly available functional element that provides security, NAT traversal, protocol repair and other functions to VoIP signaling such as SIP. A component of a Border Control Function.
SDP (Session Description Protocol)	A standard syntax contained in a signaling message to negotiate a real time media session. See RFC4566.

Term or Abbreviation (Expansion)	Definition / Description
SIP (Session Initiation Protocol)	A protocol specified by the IETF (RFC3261) that defines a method for establishing multimedia sessions over the Internet. Used as the call signaling protocol in VoIP, NENA i2 and NENA i3.
SNMP (Simple Network Management Protocol)	A protocol specified by the IETF used for managing devices on an IP network.
SRC (Session Recording Client)	The Logging Service acts as a Session Recording Server (SRS), and accepts media and metadata from a Session Recording Client (SRC).
SRTP (Secure Real-time Protocol)	An IP protocol used to securely transport media (voice, video, text) which have a real time constraint.
SSL (Secure Socket Layer)	The Secure Sockets Layer (SSL) is a computer networking protocol that manages server authentication, client authentication and encrypted communication between servers and clients.
SSO (Single Sign-On)	Single Sign-On (SSO) is a session/user authentication process that permits a user to enter one name and password in order to access multiple applications.
TDoS (Telephony Denial of Service)	Illegal attacks targeting the telephone network by generating numerous phone calls, tying up the network and preventing an agency from receiving legitimate calls.
TLD (Top-level Domain)	A top-level domain (TLD) is the last segment of the domain name. The TLD is the letters immediately following the final dot in an Internet address.
TCP (Transmission Control Protocol)	A communications protocol linking different computer platforms across networks. TCP/IP functions at the 3rd and 4th levels of the Open System Interconnection (OSI) model.
TLS (Transport Layer Security)	An Internet protocol that operates between the IP layer and TCP and provides hop-by-hop authentication, integrity protection and privacy using a negotiated cipher-suite.

Term or Abbreviation (Expansion)	Definition / Description
UDP (User Datagram Protocol)	One of several core protocols commonly used on the Internet. Used by programs on networked computers to send short messages, called datagrams, between one another. UDP is a lightweight message protocol, and compared to TCP, is stateless and more efficient at handling lots of short messages from many clients.

3 Introduction to Next Generation 9-1-1 Security

This document is intended to provide guidance on how to approach the end state security environment as documented in the i3 architecture documents. This document is not normative. The reader should refer to the actual normative standards. This document also does not directly impose any deployment requirements.

IP networks were developed to foster resilient connectivity but not security. IP multimedia services are easy targets because they are based on IP networks that are inherently insecure. IP was also developed to be flexible, so there are many types of services within today's infrastructure that have been built on top of IP over time. Once you transition to NG9-1-1 you rely on these IP networks to deliver Emergency Services.

Trying to classify and control all of the communications on your network is an increasingly difficult task, never mind understanding the content of the exchange. Circuit switched networks were designed to connect two points with a physical and unchanging path for the duration of a call. Conversely, IP was designed to route around failures and take the best path between two points. Combining the classification problem with the fact that communications are becoming increasingly ubiquitous and virtual, users do not necessarily know where one another are located or how their information is transmitted.

Overload happens when all the available resources are consumed, or there is a significant impact to the performance or availability of your network. Intentional overloads can be malicious denial of service or distributed denial of service events, which we refer to as DoS or DDoS attacks. These are conducted over the IP network, or even automated and brought in over the telephony network to tie up trunks. In the world of SIP it is most common to see invite (i.e., SIP INVITE method) or registration (i.e., SIP REGISTER method) floods from the Internet since these tie up both bandwidth and your telephony resources. DoS attacks are very difficult to stop since the attacker can shield themselves by using many servers or workstations that they have previously compromised and taken control of during a single attack. Social networks are also being used to coordinate focused attacks by willing parties.

An important design consideration in NG9-1-1 is the principal that there is "No Security in Obscurity". This concept means that we must have active, effective secure mechanisms that prevent

unauthorized access, and we do not depend on hiding access, or making resources hard to find. Attempting to achieve security through obscurity is actually quite common, but has been repeatedly shown to be ineffective as the only measure. The security issues we face arise through deliberate attempts to manipulate systems, not accidental mishaps. The level of sophistication of attackers is very high, and none of the obscurity mechanisms work against sophisticated attackers. If we can stop sophisticated attacks, we can also stop unsophisticated attackers. We put a lock on the safe, not hide it from view. Recommended practices in this document purposely depart from the concept of "Security by Obscurity" and rely on active protection.

Another important design consideration is to avoid relying on "walled gardens". Walled gardens refer to attempts to build secure boundaries limiting access to a network, and then assume that whatever is inside the network is safe. Creating secure borders to networks is certainly a primary defense mechanism and is highly encouraged. However, it has proven nearly impossible to maintain the borders securely, and thus such borders are regularly breached, not only by intruders, but by well-meaning insiders trying to get their job done. The need to well defend the interior of the NGCS at all layers is further necessitated by successful attacks which managed to breach one system from the outside, and use that breach to exploit vulnerabilities of other systems from the inside. Thus we recommend that the inside of the NGCS be treated as if it was the open Internet, and all activities within it be protected. The standards require all external interfaces on all functional elements use secure protocols, and make use of uniform identity, authentication, authorization, privacy, integrity and non-repudiation mechanisms. Even systems beyond the NG9-1-1 standards should employ similar security mechanisms to protect them from harm.

NENA-STA-010 includes a mechanism to control access to data ("data rights management" – DRM). Data that might be considered sensitive should be subject to data rights management. This means all vendors must implement the mechanisms wherever data is created or used, and authorities deploying such systems must carefully consider how the mechanisms can best be used to minimize threats while allowing agents to do their jobs. The latter involves careful consideration of agent roles, and construction of the DRM rule sets, secure credential handling and appropriate handling of errors and alarms.

In extreme circumstances, a PSAP may need to transfer operations to another facility. Security considerations for such a transfer must be evaluated well in advance, to assure that the alternate facility can securely function as a replacement for the original. In some circumstances, the roots of trust (certificates and certificate authorities) for the primary and backup facilities could be different, and consideration of such situations must be evaluated for their effect on agents and agencies that may need to communicate with the backup facility.

3.1 Request for Information and Compliance

When developing a request for proposals or information, security considerations must figure prominently. Writers of such requests should carefully review NENA 75-001 and 75-502. Every element in a proposed bidder response should be required to follow all the security requirements in NENA-STA-010 for that element. For baseline requirements within the request to bidders, use the

following table. Also, there is a blank version of this table in Appendix A for use with a security checklist.

	Security Element								
Functional Element	TLS	SRTP	Single Sign on	Credentials traceable to PCA	Patches	Communicate with policy mechanisms	Call path	Databases	External (to ESInet) interface
LNG	√	√	Note 1	Note 1	√	√	√	n/a	√
BCF	√	√	√	√	√	√	√	n/a	n/a
ESRP	√	n/a	√	√	√	√	√	n/a	n/a
ECRF	√	n/a	√	√	√	√	n/a	√	√
LVF	√	n/a	√	√	√	√	n/a	√	√
LPG	√	√	√	√	√	√	√	n/a	√
Call Handling	√	√	√	√	√	√	√	n/a	n/a
Dispatch	√	n/a	√	√	√	√	n/a	n/a	n/a
IDE	√	n/a	√	√	√	√	n/a	n/a	√
Outgoing Alert	√	n/a	√	√	√	√	n/a	n/a	n/a
Incident Handling	√	n/a	√	√	√	√	n/a	n/a	n/a
IMR	√	√	√	√	√	√	√	n/a	n/a
Map database	√	n/a	√	√	√	√	n/a	√	n/a
MIS	√	n/a	√	√	√	√	n/a	√	n/a
RMS	√	n/a	√	√	√	√	n/a	√	n/a
Mobile data	√	n/a	√	√	√	√	n/a	n/a	n/a
Logging Service	√	n/a	√	√	√	√	n/a	√	n/a
Media recording interface	√	√	√	√	√	√	√	√	n/a
Radio interface	√	√	√	√	√	√	√	n/a	n/a
Bridge	√	√	√	√	√	√	√	n/a	n/a
Agency locator service	√	n/a	√	√	√	√	n/a	√	n/a
ADR	√	n/a	n/a	n/a	√	√	n/a	√	√
Note 1: If LNG is provided by 9-1-1 authority then SSO and credentials traceable to PCA apply.									
Note 2: Nearly all FE's will need to enforce data rights management									

Security Elements

This section describes how purchasers might evaluate responses to deploying NG 9-1-1 systems and security questions in Request For Information/Proposal/Quote (RFI/RFP/RFQ) processes. Deploying NG 9-1-1 includes deployment of an IP transport infrastructure which may be provided by an existing Internet Service Provider (ISP). Such ISPs would have their own security infrastructure which may not conform and should not be expected to conform to NG9-1-1 specific security mechanisms. In evaluating ISP security mechanisms, all of the questions should be asked, yet the responses need to be evaluated without reference to specific NG9-1-1 mechanisms.

1. TLS
    a. What version of TLS do you support? For example, TLS Version 1.2, RFC5246 is the current standard and should be implemented by all elements.

b. How are credentials managed? Management of private keys for TLS in a secure manner is not easy without some kind of tool. Consider the entire "chain of custody" for the private key and evaluate where the weaknesses are. Ask about how state in the server is backed up – is the key safe? Also ask about how private keys are managed where there are redundant elements. If a certificate is issued for an element named esrp.state.pa.us, but there are 3 instances of the server that provides the function, there would need to be three copies of the key. How are the copies maintained in a secure manner? What happens if credentials are compromised? What steps would need to be taken to re-key, and how long would that take? How are expired keys managed? How do you make sure new keys are created and installed before the old keys expire?
c. Are persistent TCP/TLS connections supported and how are they managed? Persistent TCP connections are not needed for security; they are a tool that can be used to mitigate the time it takes to negotiate a new TLS connection. If a TLS connection is maintained even if there is no current traffic, it can be reused when a new transaction is started, rather than negotiating a new connection. So, it's good to have that capability. Persistent connections could be created automatically assuming that a transaction between two elements is a good predictor of further transactions in the near future. This would incur a TLS connection for the first one, but subsequent transactions between the elements within the time-out of the persistent connections would be fast. They have the advantage that no work is needed. Another strategy is to provision them where you know in advance that transactions between specific elements are common. This has the advantage that even the first transaction can use the connection (which would be established at boot or restart) and does not require a timeout mechanism. The downside is that provisioning is needed, and predicting where connections are frequent isn't necessarily easy. Persistent connections consume some resources, so you don't want to maintain a lot of them for long periods of no use.
d. What cypher suites are supported? NENA-STA-010.2 specified a minimum of RSA-1024, AES and SHA-256. Support for more aggressive ciphers should be encouraged, for example, elliptic curve, but recognize that a "bid-down" attack might be possible since to meet the NENA-STA-010 standard, the minimum cipher suites must be supported.
e. What bid-down prevention mechanisms are available? Systems internal to the ESInet should not allow less than the minimum cipher specified in NENA-STA-010. Connections to external networks may need to accept less aggressive ciphers. There is always a tension between lowering security and getting 9-1-1 calls through no matter what. Generally speaking, the latter ought to prevail, but perhaps network operations and call takers could be informed when security was compromised in favor of getting a 9-1-1 call through.
f. What is the source of your TLS stack? There are some less tested and less secure implementations. Be aware that even highly regarded stacks, such as OpenSSL

have been found to have vulnerabilities, but stacks that are not well known and tested have fared worse in most real world attacks.

g. What happens if a request is made without TLS? As with a less-than-minimum cipher suite, allowing a 9-1-1 call without TLS must be acceptable, but notification and follow up with the source would be useful to avoid future problems.
h. What alarming and reporting do you support relative to TLS and connections? Failures of security mechanisms in NG9-1-1 elements should always result in some form of follow up. All forms of alarms have the issue that too many of them can obscure the underlying cause. It is important to understand how failures are handled so that problems can be addressed rapidly and accurately so that service is always available while maintaining secure operations.
i. Does your stated capacity assume all connections use TLS/TCP? It is often the case that performance specifications do not include establishment of a TLS connection. While, as above, persistent TCP connections can ameliorate the time penalty of establishing TLS, performance when TLS must be negotiated should always be the base performance standard, with the optimization methods as secondary specifications.

2. SRTP
   a. Does your stated capacity assume all calls use SRTP? As above, it is common to quote performance without considering SRTP overhead. Purchasers must understand if their systems can support their performance requirements when all security mechanisms are employed.
   b. What keying mechanisms do you support?
   c. What cypher suites are supported? As with TLS and SIP, the same issues of cipher selection apply to SRTP.
   d. What media types are supported by your SRTP implementation? Carefully consider the effect of SRTP on video calls. Many systems were optimized for voice and cannot support SRTP on video. While text (RFC4103 Real Time Text) is low bandwidth, it too needs to be protected via SRTP
3. Single Sign-On
   a. Do you support the NENA-STA-010 Single Sign-On mechanism? The only required function in NENA-STA-010 that requires the SSO mechanism is the data rights management function, which can use agent identity and role in authorization decisions. However, if there is any other function that controls access by individual user, it should use the Single Sign-On mechanism to provide authentication of the agents.
   b. Is there any other user authentication mechanism in the system? In general, if no other user authentication mechanisms are allowed, the SSO mechanism should be used. The security concern is how many usernames and passwords users need to maintain. Research shows that if a user has to handle many different username/password combinations, they exhibit bad behavior.

c. What are the mechanisms used to support authorization decisions? The standardized mechanism is the NENA-STA-010 mechanism that uses a technology called "XACML" which is well understood. Other mechanisms may not have the flexibility to allow fine-grained authorization decisions.
d. Where can authorization policies be stored? NENA-STA-010 defines a Policy Store to hold all kinds of policies. It is desirable to maintain policies in a small number of well controlled places. The NENA-STA-010 mechanism uses its own data rights management system to control access to policies.

4. Credentials traceable to PCA
Every Agency and every Agent in that Agency needs credentials. Credentials come from a Certificate Authority, and in NG9-1-1, a PSAP Credential Authority. There will be, we hope, a National PCA soon, but in most cases, Agency credentials should come from a state-level PCA, where the state PCA gets its credential signed by the national PCA. Agents should get credentials from a regional PCA, whose credential would be signed by the state PCA. The level of concern for an agency credential is necessarily higher than an individual agent, and an agent with high authorization has more concern than an agent with minimal authorization, but all credential maintenance processes should consider:
   a. What are the limits on credentials? You should consider expiration times and stated use limits the CA allows.
   b. How are credentials manipulated? What is the "chain of custody" for the private key and the X.509 certificate containing the public key?
   c. How are credentials stored? What vulnerabilities exist for credentials?
   d. How are credentials added, deleted, updated? What processes are used? How are they evaluated for vulnerabilities? What mechanisms exist to handle compromise of credentials? How are expiration of credentials handled to keep service seamless?

5. Patches
Every computer system is vulnerable to attack. New vulnerabilities are discovered constantly. Keeping systems up to date with security patches is a vital process that must be evaluated for every computer-based system, even embedded systems. While some new vulnerabilities are only theoretic and do not pose imminent threat, others may require immediate action. Understanding how each system is updated is vital to maintaining secure environments.
   a. How are newly discovered security vulnerabilities reported? Your IT security team must receive timely notification of new vulnerabilities. If you are reading about a new vulnerability in a trade magazine, your vendor should have a fix in process with a short deployment schedule, but getting notice from your vendor to you is the first step.
   b. How are newly discovered security vulnerabilities resolved? It is usually not acceptable to wait for the next feature release. Some form of patching is needed.
   c. What is the typical and maximum time between reporting vulnerability to resolution? A high priority patch should take days to resolve, while other

theoretic-only vulnerabilities might be repaired in times measured in weeks or even months.

d. What mechanisms do you use to discover and learn about vulnerabilities in third party and open source code used within your product? Vulnerabilities are most often encountered in subsystems and libraries used by many organizations. It is often difficult for vendors to recognize when their code is vulnerable. Understanding how they monitor patch notices from their component suppliers is useful to evaluating these vendors.

e. What software components are used in your solution that are not the most current version available and are patches or updates being supplied for new vulnerabilities on the older versions? Vendors sometimes delay or avoid new versions of software they incorporate in their solution, sometimes because the new version requires changes to their code and they don't want to or have the resources to make such changes. Running on an older version may result in unpatched vulnerabilities. Sometimes the supplier of the code does make patches available for older versions, at least for some time, so it is not necessary that every component always run the most recent version. Purchasers must understand where their suppliers are not up-to-date, and must be assured that vulnerabilities discovered will be addressed promptly.

f. What mechanisms do you use to deliver a patch? Once a patch is available, it has to be installed on all vulnerable systems. How code is patched, what expectations for downtime are needed and what actions your staff will need to complete are important to understand.

6. Communication with policy mechanisms. NENA-STA-010 has a uniform policy mechanism based on XACML. All policies should use this mechanism.
   a. What NENA-STA-010 policy mechanisms are supported in your product?
   b. What additional policy mechanisms are supported? Implementations may have additional policy controlled processes. Ideally, these would use the XACML mechanism.
   c. Where can policies be stored? NENA-STA-010 describes a policy store with a specified interface. A system may have an internal policy store, but the system should be able to use an external policy store with the standardized interface, and the internal policy store should have the specified provisioning interface in order to use common tools to edit policies and to control access to resources in a uniform, controlled manner.
   d. If your system does not use the mechanisms specified within NENA-STA-010, what encryption mechanisms does your policy interface support? Where the standardized mechanisms are not supported, it is important that policies be securely stored, and mechanisms that control who can modify policies be secure.
7. Call path
   a. Does your stated capacity assume all calls use SIPS/TLS, SRTP? It is all too common to specify call-handling capacity assuming no TLS and/or SRTP. Security mechanisms involve cryptography that is compute-intensive. You need

to be able to maintain your requirements on call processing when all security mechanisms are actively employed on all calls.

b. What happens if a request is received without SIPS/TLS, SRTP? While use of the security mechanisms is highly desirable, incoming calls should not fail if security mechanisms are not offered or the crypto fails. Such calls may be more suspicious than calls that successfully employ the mechanisms. Note that SIPS/TLS and SRTP are hop-by-hop mechanisms. Thus even if a call is received that does not use them, all hops inside the ESInet should use them.

c. If your system acts as Session Recording Client (SRC) how does it handle keying on the recording session? Maintaining secure media includes maintaining the security of the logger, and the connection from the source of the recording to the logger.

8. Databases

a. What security mechanisms are available on your system's databases? The databases described by the NENA i3 specifications include the ECRF, LVF, the data source for the MCS, GCS and Map Display Database, the policy store and the Agency Locator. Also described, but external to the ESInet is the LIS, CIDB, ACDR and ALDR. Many systems have internal databases. Each of these databases should have controls on who can access or modify the content, ideally using the standardized access control mechanism, but at least using the same identities and roles. It is difficult to securely maintain parallel user identification, authentication and authorization mechanisms. Sometimes, implementations use internal identities on the underlying database and use the standard mechanism on the system functions that actually manipulate the database. If correctly implemented, and if the internal mechanisms and credentials are secure, such approaches can be secure. If there is a single identity, and the password for that identity is stored in plaintext in some configuration file (an all-too often encountered circumstance), security is probably inadequate.

b. What type of authorizations for add/modify/delete/read are available? There is always a tradeoff on how fine grained the mechanisms offer and how difficult it is to use them in a way that maintains security. Often the response to more complex systems is to escalate permissions to get around problems. Sometimes a simple system won't let management properly segment permissions. A balance is needed.

c. What forms of credentials are available? This is back to the overall credential issue. If the database has its own user ids and control mechanisms, how are they protected, and what form of information is needed to access the database. As before, username and password is not considered acceptable any longer, although if the database is not accessible, in any way, from the Internet, and the systems it is installed on have adequate security, and the mechanisms used by the database to protect the data are adequate, there may not be a need to insist on 2 factor authentication to the database itself.

d. Is the database encrypted? Some databases used in 9-1-1 contain data that must be kept private. Other information is public. Where data must be kept private, the database where it is stored must usually encrypt the data to prevent unauthorized access.
    i. If not, what protections are used to prevent unauthorized access or modifications? It is conceivable that there is no way anyone can actually access the storage mechanism used by the database which may obviate the need for encryption. Even in such circumstances, encryption of the database is recommended.
    ii. If so, how is the key protected? While most commercial databases have secure mechanisms to maintain keys, it is surprising how often the applications do not adequately protect the keys.

9. External (to ESInet) interfaces. Wherever the ESInet is connected to another network, the possibility exists that that network may be compromised, or open to attack, directly or indirectly, which can then be used to attack the ESInet. Therefore, every connection to another network requires some way to protect against attack. Primarily, this is provided by the BCF. "Network" should be understood in a broad context. So, for example, an LNG connected to a legacy origination network is subject to attack and thus the LNG that connects it to the ESInet must be protected by a BCF (and therefore outside the ESInet). There should be no exceptions to this rule. There are no such things as "trusted" networks although ESInets can generally trust each other to place BCFs in front of all sources of packets and calls from external networks.

   Mitigation must consider both the available bandwidth on the interfaces as well as the packet processing capability on each interface. It is not unusual for specialized mitigation services to be engaged when attacks are detected. Such services would use changes in routing announcements (BGP) to direct traffic to the mitigation service, and then a private connection is used to direct valid traffic back to the ESInet. The mitigation not only has to handle the largest size attack in raw gigabits of bandwidth, but must also withstand the maximum packet rate on those interfaces (lots of small packets)
    a. What DDoS protection is available? What is its capacity? As of the date of release of this document, the largest observed attacks were in excess of 600G. If your network cannot withstand an attack that large, then your network can be overwhelmed. DDoS mitigation of 600 G attacks is available. There are no real excuses why an ESInet cannot be protected against the current largest attack. Attack and mitigation technology is an "arms race" and it's always the case that for some period of time the attackers may have more capacity than the mitigation, but that should be transient. Mitigation against generic packet attacks, DNS attacks and SIP attacks are required.
    b. What ports are open normally, and how controllable are they? This is a BCF configuration issues. Where packets for specific ports are allowed into the network, mitigation for attacks against those services is required. Obviously,

minimizing which ports are open is a good strategy, but agencies have to be able to do their jobs, which require many common services to be available.

c. Describe how attacks are detected, and what resources are available from your company to mitigate? Attack mitigation is a specialized discipline. The processes, skills available, escalation mechanisms and other aspects of mitigation must be clearly understood, and deemed adequate. Agencies may need to engage consultants to evaluate how prepared their vendors are for attack.
d. What happens to the system under attack? Consider both external attack and internal attack. These days best practice is to use DDoS mitigation services to enhance the ability of systems to withstand very large attacks, but even the primary path must be sized to withstand large scale TDoS attacks that may not be amenable to mitigation from such services. Ask vendors
    i. What are the measures that you use to ensure resiliency under attack?
    ii. Can you provide examples of attacks for which your system was tested?
    iii. What types of overload/overflow can be transmitted to your interfaces?
e. What known attack types can your system mitigate?
f. What services outside of your system does your product rely on to maintain service even while under attack?
g. Are your external interfaces physically and logically separate from your management interfaces? If attacks on the system can overwhelm the management interfaces, it may not be possible to manage the network under attack. Strict separation is required, with separate, protected paths.

10. Security and site failover. If a single point of failure exists, an attack may target that SPOF, thus taking down the service, or another failure in the system may render a security mechanism inoperable. It is essential that no SPOF exist in the network for 9-1-1 critical services.
    a. What single points of failure might exist within the security architecture?
        i. For example, the database which stores credentials.
    b. How is security maintained across diverse systems and failures? Refer to NENA-STA-010 section 3.8.
    c. Security failures cannot compromise the overall function of the NG9-1-1 system. For example, how does your Firewall handle site failover? As per NENA NG-SEC 75-001, Use of redundancy and/or diversity can have an effect on various types of security products. Most notably, traffic failover between different cities and different firewall sites can result in dropping sessions which are underway at the time of the failover.
    d. What level of availability are you providing in terms of 99.xxx% and will the degrees of cost vary based upon availability?

3.2 Cryptographic Mechanisms

To protect communications, we use cryptography. There are four unique mechanisms required to establish and maintain secure communications:

- Authentication, which identifies the identity of the parties that are communicating
- Authorization: which identifies what operations an authenticated party can do
- Secrecy: which protects communication from unauthorized viewing
- Integrity checking: which protects communications from being modified in transit

We also sometimes need "non-repudiation" which allows a receiving party to know that an operation was completed by its correspondent in a way that does not permit the correspondent to subsequently claim it did not complete the operation.

In NG9-1-1, we use:

- RSA based Public Key Cryptography using X.509 certificates to authenticate elements, agencies and agents
- A XACML based Data Rights Management system to control authorization
- AES based Encryption to provide Secrecy
- SHA based Digest based digital hashing to provide integrity protection
- Dsig based Digital Signatures to provide, among other things, non-repudiation

3.3 Certificate Management

3.3.1 What are credentials and who needs them

Every Agency, Agent and every "addressable" entity (meaning an element that is in the call path, or is in the path of any web service, or any other NG9-1-1 protocol) needs unique credentials. NG9-1-1 standards prohibit sharing of credentials. No two agents should use the same credentials, including service technicians. Agency credentials are rarely used within the ESInet, but an agency may issue a credential to an element, signing the element's credential with the agency credential. Each element and each agent of an agency would use their individual credentials.

Identity is the mechanism to identify "who" is doing something. In NG9-1-1, a number of entities have identities:

- Agencies (organizations such as PSAPs, Fire Departments, state 9-1-1 Authorities, as well as vendors and other enterprises who supply services to the ESInet)
- Agents (people who work for agencies)
- Elements (computer systems or other automata who perform useful functions in the ESInet or within an agency's local network)

We issue credentials to identities. These credentials have several components:

- An identifier, which is unique, and is used throughout the system
- A set of authentication factors, such as passwords, fingerprint scans, tokens etc. used to prove that an entity is who they say they are

- A public/private key pair that uses Public Key Cryptography for protocol authentication mechanisms
- An X.509 Certificate that contains the identity, the public key, and a signature of a recognized Certificate Authority attesting to the authenticity of the information in the certificate

Agents are authenticated to the system using a standard "single sign-on" (SSO) mechanism. This mechanism uses user ids, passwords and biometric data to access the private key assigned to the agent. There is a very limited set of services, such as SNMP that cannot use NG9-1-1 specified single sign-on mechanism. Vendors should be asked which systems cannot support the SSO mechanisms and requestors should minimize exceptions to the SSO requirement. Security for all services, including SNMP should always be required, even if the SSO mechanism cannot be used.

Vendor technicians are an important class of "agent" that must be carefully considered. Each must have their own credentials, no sharing should be permitted. While it is possible that the contracting agency issues the credential used by a technician while working on behalf of that agency, it is problematic for a technician who supports multiple agencies to have to maintain separate credentials for each one, especially if passwords are one of the authentication factors. It may be preferable for vendors to obtain an agency credential, and then issue credentials to its agents. This would require the contracting agency to permit access to its systems by agents of the vendor agency.

3.3.2 Certificates and Authorities

X.509 certificates are the method by which public keys are distributed. The certificate is issued by a Certificate Authority (CA). The CA collects information from the agency or agent requesting a certificate sufficient to prove that the requestor is authentic, and then creates the certificate, signing it with its private key.

CAs are arranged in a tree, with a root CA the so called "root of trust". The certificates of the intermediate CAs are signed by the next level up in the hierarchy. The root certificate, which is self-signed by the root, is widely distributed. By examining a cert, a relying party can check the signature of the CA to validate that the cert is genuine. It can validate that the signature of the CA's cert (which it would use to get the public key of the CA) is valid by obtaining the public key of the next level up CA, and validating it. It can follow the chain all the way back to the root CA, assuring that the certificate it started with is valid. The tree of CAs is called a "Public Key Infrastructure" and is needed to effectively use a public key based cryptosystem like that specified for NG9-1-1. Certificate Authorities operate under two important documents, a "Certificate Policy" and a "Certificate Practice Statement". Establishing a CA is much more than a technical capability – the processes by which certificates are created, maintained, revoked and replaced must be adequate to the task. The CP/CPS defines the practice and the methods used by the CA to meet the requirements.

3.3.3 Establishing a National PCA

The National PSAP Certificate Authority (PCA) **should** exist, and would be the root CA for the NG9-1-1 PKI. Its cert should be cross signed by the Federal Bridge CA, which signs root certs for all US federal government CAs as well as many other national CAs. This allows "federation" of PKIs. For example, police departments could create their own tree of CAs, with their own root CA. If that root CA's cert was signed by the federal bridge CA, then a PSAP with a cert signed by the National PCA could mutually validate with a police department ("mutually validate" means both agencies can check the cert of the other). NENA has recently agreed to establish the national PCA.

The National PCA should establish a process relative to how States establish the notion of a local CA and how individual agents get a CA. If States begin deployments without having established a National PCA there will be (interoperability) issues with not having a trust root. The FCC, Program, DOT, APCO, and Canadian authorities are aware of the issue. Consideration needs to be taken for trust levels, transit of trust, trust tree, and who will do this. Consideration must be taken for International interoperability/capability.

A business case and technical risk analysis is required to determine the scope, resources and schedule along with procedures and staff required. This potentially could be outsourced and thus may need to be run through RFx process to build a budget.

3.3.4 Establishing a State PCA

In most cases, state PCAs should exist, and should have their certs signed by the national PCA. This will be the second level. In some states, the state and/or regions may form a statewide CA via interlocal agreement and in other states regional PCAs will exist, and will have their certs signed by the state PCA. Agencies will need to operate their own CA, or the next level up CA could provide agent certs for the agency. This would create a 2-4 level CA hierarchy (3 levels would be national -> state -> regional -> PSAP; two levels might be national and state).

States should have a State or statewide PCA. The 9-1-1 Authority and/or a group of regional 9-1-1 authorities via interlocal agreement may provide this function, or contract with a qualified contractor to provide it. Another state agency may have the capabilities to provide the state PCA on behalf of the state 9-1-1 Authority.

Regional PCAs should be considered, if needed, whenever there are adequate resources available to provide one and be responsible for it. If there are regional PCAs, they are much more likely to have simple mechanisms for verifying requests from agencies for certificates. Furthermore in situations where the local agencies cannot reasonably provide agent certificates, regional PCAs will be much more convenient for agents to provide authentication factor information than state agencies could. Like state PCAs, using a qualified contractor to provide the regional PCA service may be appropriate.

Individual PSAPs would likely only approach the PCA issue, if needed, by contracting with a qualified vendor. Otherwise, the regional or state PCA should provide certs for PSAP agents. A PSAP PCA's only task should be to issue its agents and elements their credentials. While it is

convenient to have the PCA be local, if its security and responsibilities cannot adequately be assured, then local PCAs should not be provided.

3.3.5 PCA Policies

A certificate policy focuses on certificates and the CA's responsibilities regarding these certificates. It is a set of rules that determine if a certificate is applicable to either a community or a class of applications that have common security requirements. It defines certificate characteristics such as usage, enrollment and issuance procedures, as well as liability issues.

The following references define a certificate policy

A certificate policy typically answers the question about what purposes the certificate serves, and under which policies and procedures the certificate has been issued. A certificate policy typically addresses:

- How users are authenticated during certificate enrollment
- Legal issues, such as liability, which might arise if the CA becomes compromised or is used for something other than its intended purpose
- The intended purpose of the certificate
- Private key management requirements, such as storage on smart cards or other hardware devices
- Whether the private key can be exported or archived
- Requirements for users of the certificates, including what users must do if their private keys are lost or compromised
- Requirements for certificate enrollment, renewal and revocation with provisions for an expedited revocation process
- Minimum length for the public key and private key pairs

The certificate practice statement (CPS) translates certificate policies into operational procedures. The certificate policy focuses on a certificate; the CPS focuses on a CA. It defines the way that a CA issues certificates.

A CPS might include:

- Positive identification of the CA, including the CA name, server name, and Domain Name System (DNS) address
- Certificate policies that are implemented by the CA and the certificate types that are issued
- Policies, procedures, and processes for issuing, renewing, and recovering certificates
- Cryptographic algorithms, cryptographic service providers (CSPs), and the key length that is used for the CA certificate

- Physical, network, and procedural security for the CA
- The certificate lifetime of each certificate that is issued by the CA
- Policies for revoking certificates, including conditions for certificate revocation, such as employee termination and misuse of security privileges
- Policies for certificate revocation lists (CRLs), including where to locate CRL distribution points and how often CRLs are published
- A policy for renewing the CA's certificate before it expires

RFC 3647 provides a detailed description of what a CP and CPS must contain. All CAs in the PCA hierarchy must have CP/CPSs that conform to RFC 3647.

It is required that private keys for national, state and regional PCAs be kept in "Hardware Security Modules" with FIPS-140 Level 3 validation. The same is recommended for PSAP PCAs, but Level 2 validation is acceptable, and in some cases, even Level 1 might be appropriate.

All PCAs, regardless of level must have an independent audit examining their actual performance against their CP/CPS. If a higher level PCA is issuing a cert to a lower level PCA, not only should it be especially vigilant in assuring that the request for a cert is valid, but it must also review the CP/CPS, and audits of same to assure itself that providing a cert to that PCA is within its own CP/CPS.

The higher in the hierarchy, the more stringent the issuing PCA should be in determining the validity of the request for a certificate. The National PCA should consider in-person visits to state PCAs with an examination of equipment and processes before considering issuing a cert. State PCAs may need similar qualification of regional PCAs. If there is no regional PCA, some secure method of verifying the identity of a request by a PSAP for a cert must be provided. Entities issuing certificates for agents should require agents to present themselves in person (to gather the required authentication factor information). Some secure mechanism to assure that an agent candidate is authorized by the agency must be provided together with something that links the agency authorization with the candidate presenting him or herself to the PCA.

3.3.6 Credentials for entities outside the ESInet

Agencies other than PSAPs may need certificates. For example, a contractor operating an ESInet on behalf of a 9-1-1 Authority should have its own agency certificate unless the contracting agency is willing to assume the liability of issuing certificates to the contractor for its elements and agents. State or regional PCAs would issue such agency certificates. As with PSAPs, whether such agencies run their own CAs to issue certs to agents and elements depends on the adequacy of the controls the agency can provide to managing the keys and processes for the CA. If the agency does not possess the requisite skills, it should be required to use a regional or state PCA to issue its agent/element certs, or it should contract with a vendor who has such skills.

In the ESInet environment, when calls are transferred from a PSAP to non-ESInet entities, i.e. tow truck company, secure communication is still required for delivery of critical information including

EIDDs. The external entities must have credentials, ideally issued by an external CA trusted by the internal agencies that must interact with them. In some circumstances, outside entities may need to obtain credentials issued by a PCA. Systems that communicate with outside entities will need to be provisioned with the root CA certificates of any external PKIs.

3.4 Authentication

Authentication is the process of securely identifying an entity. For agencies and elements, typically RSA or elliptic curve algorithms allow two entities to mutually authenticate each other. Mutual authentication is mandatory for nearly all NG9-1-1 interactions.

3.4.1 Single Sign-On

NENA standards specify that a single sign-on mechanism using a specified SAML based mechanism must be used by all services requiring agent authentication or authorization in NG9-1-1. The credentials issued by a PCA are the credentials used to identify agents – the SSO provides access to the private keys, and provides a standard mechanism to authorize activities undertaken by agents. No other mechanisms for identity should be used.

For agents, the SSO mechanism is used to authenticate the user with the SSO. The SAML protocol can then be used to provide other services with authenticated interaction with users. To authenticate a user using SSO, the user must provide at least two "factors", one of which is typically a password or pass phrase. For the second factor, biometrics is recommended, but other mechanisms may be used.

1. Use of hard tokens can be expensive since users tend to lose them. Consider whether a hard token such as a USB key, smart card, or sequence based token with an LCD (such as a SecurID) is right for you. Soft tokens using apps on mobile phones are popular. If the phone has a biometric access mechanism, consider requiring it.
2. Another alternative is use of SMS pin messaging to mobile phones. While this may seem attractive at first, consider that the availability of the mobile network in a disaster situation will be limited or non-existent.
3. If a biometric identifier is chosen, fingerprint readers are the most reliable. However keep in mind that unless the reader and communication to/from the reader is protected by encryption that they may be susceptible to replay attacks.
4. When selecting equipment that will integrate with your two factor auth deployment, look for support for RADIUS, TACACS+, and Active Directory for networking gear.
5. While alternative mechanisms like OpenID are attractive, they do not replace the NENA-STA-010 PKI and SSO mechanisms. OAuth could be useful.

3.5 Deploying TLS

NENA-STA-010 states: "All protocol operations must be integrity-protected (via TLS or IPsec)."

Although the majority of existing SIP deployments is transported over UDP and thus not secure, 9-1-1 is not an ordinary SIP deployment and all SIP communications in the ESInet must be secure.

The IETF standards on which i3 are based require TLS end to end. To the extent possible, 9-1-1 Authorities should work with origination networks to implement end to end security.

Transport Layer Security (TLS) is a security protocol that encapsulates other application protocols. TLS is an IETF standard based on Secure Socket Layer (SSL). TLS is used for confidentiality, integrity and data compression of SIP signaling only, not media. Using TLS prevents eavesdropping, tampering, and forgery. TLS can be used with any protocol that uses TCP, which includes HTTP (used by i3 web services) and SIP, used by calling services. The current standard goes to version 1.2. TLS carries SIP over TCP only after TLS layer established.

NENA standards require "mutual authentication" using at least RSA-1024, which implies both ends have an X.509 certificate available to the other party. See above for how certs should be issued.

Certificates issued to redundant components (especially those that are involved in site to site geographic failover) will need a certificate with one common name and one or more alternative names. These "names" are the DNS fully qualified domain names of the component at its respective site. Wildcard certificates can certainly be issued in these situations as well, but they will be less secure and are not recommended

Use 2048 bit certificates. Not many CAs even issue or recommend 1024 anymore.

Certificates need to be verified; there are circumstances where a certificate is revoked before it expires. To verify a certificate, there are two mechanisms. CAs publish a "Certificate Revocation List" (CRL) which can be downloaded and checked against a cert. A protocol known as "Online Certificate Status Protocol" (OCSP) can contain the CA and ask if the cert is still valid in real time. Either mechanism can be used, although OCSP has the advantage of doing real time verification. PCAs should be implementing both mechanisms.

Certificates have an expiration date, and they cannot be used past that date. Typically, a new cert is issued before the old one expires. Managing certificate expiration is a process that all elements in an ESInet have to handle, and understanding what is required is essential to maintaining security. It is much too common to see certificates expiring without a new one installed, creating an insecure environment.

Under NO circumstances should certificates be self-signed (except for the root CA). No user or administrator should ever see, let alone agree, to a security warning that a certificate is invalid or the authenticity cannot be verified. Careful management of element names and the certificates for those elements is required to avoid a situation where an element is reached with a URL that is not one in the certificate's identity field.

There is an alternative to TLS: IPsec. IPsec creates a secure tunnel between two elements or two networks. The cryptographic functions used with IPsec are the same as those used with TLS, and thus an IPsec tunnel with the appropriate crypto-suite is as secure as a TLS connection with the same crypto-suite.

There are two problems with IPsec as used in an ESInet: the communicating elements typically have no way to know if the IPsec tunnel is operating, and the mechanisms are more susceptible to failure in major disruptions, because all alternate paths have to be considered in advance and IPsec tunnels

arranged, or pre-configured. TLS can be established as long as an IP path exists between the communicating entities. For this reason, IPsec is discouraged, but permitted.

It is not uncommon to find that IPsec tunnels are used by network engineers to build networks, and then application engineering deploys TLS within the tunnel. As before, the problem with IPsec is there is no way for the communicating elements to know that IPsec is established, so to be sure, TLS is deployed. This problem has no easy solution, so in fact TLS really is preferred.

The often-cited advantage of IPsec is that it is established when routers are booted, or the elements that the tunnel runs between are initialized, and there is no time/compute penalty at the beginning of a session (call, web service connection, etc.) that TLS establishment often encounters. To ameliorate this effect, implementations can make use of "persistent TCP" connections. These are TCP connections either established at boot time, or maintained beyond the first use of the connection between two elements. When a new connection is needed, the persistent connection can be reused, avoiding the crypto setup time. ESInet elements should be able to use persistent TCP to achieve rapid response. Purchasers should understand how persistent TCP connections are established and maintained. They can be completely automatic, where as soon as a connection is established, it is maintained persistent for some configurable time, after which it is allowed to drop. Where there are relatively frequent connections between two elements, there is likely to be a persistent connection available. The downside of this automatic mechanism is that the first connection between two elements always encounters a delay to establish TLS, and if there is an extended period of time with no use, the connection drops and has to be established again when the next transaction occurs.

Alternatively, persistent TCP can be provisioned. The agency decides in advance where it wants persistent connections, it provisions the pairs of elements in the systems and when the elements boot, the connections are established, and maintained permanently. The downside of provisioned persistent TCP is that effort must be made to predict where such connections are needed, the provisioning must be completed accurately, and if there are circumstances where there in fact are frequent interactions between a pair of elements that is not predicted, there will be TLS startup delays for each connection.

Probably the best way to implement persistent TCP is to allow both mechanisms, so that first transaction delays are eliminated for known-frequent pairs of elements, but automatic mechanisms keep other connections persistent as events dictate.

TLS is a "hop-by-hop" mechanism, as opposed to an "end-to-end" mechanism. A SIP call my traverse several SIP Proxy servers, including one or more ESRPs, plus one or more B2BUAs, including BCFs in the path from the calling device to the call taker workstation. TLS must be deployed on each of the "hops" along the path to create a secure connection. There are no practical ways of providing end to end security in SIP (the standards describe using S/MIME, but implementation of S/MIME has proven to be impossibly difficult for large scale SIP systems.

TLS encryption will be more difficult to troubleshoot; i.e., test equipment may not be able to decrypt and decipher a problem in a call flow. The network availability could be affected by time to repair.

3.6 Deploying Secure Real-Time Transport Protocol

The Secure Real-Time Transport Protocol (SRTP) provides encryption and authentication media path (TLS protects the SIP signaling path). RTP and RTCP traffic are encrypted as described in RFC3711, The Secure Real-time Transport Protocol (SRTP). The negotiation and establishment of keys and other cryptographic materials that support SRTP is described in RFC4568, Session Description Protocol (SDP) Security Description for Media Streams. Cryptographic parameters are established with only a single message or in single round-trip exchange using the offer/answer model defined in RFC 3264, An Offer/Answer Model with the Session Description Protocol.

Interfaces should be encrypted (TLS/SRTP or IPsec) to prevent attacks such as Man in the Middle (MITM). External and internal interfaces are vulnerable. Consider potential MITM on internal networks and sensitivity of communications.

A "Man in the middle" attack like session-hijacking, media injection, and eavesdropping generally relies on some sort of physical presence on site to connect to the network, so they are more difficult and are not as frequently observed. You can take technical precautions like authenticating devices that connect to your network and disabling or firewalling connections in public areas to keep them from connecting to the phone system.

This is where you have to ask yourself how good your physical security and employee security training programs are, or what level of trust you give your user population. Be aware that a short visit to your building or a compromise to your wireless network might enable an attacker to install either hardware or software that will enable future man in the middle attacks from a remote network.

While local attacks are easier, you can't neglect the fact that you'll see remote attacks too. Don't discount the role that endpoint security products fulfill in addition to your employee training. Man in the middle attacks are becoming easier to perform remotely using tools like the Social Engineering Toolkit, Metasploit, and Armitage. These tools make it easy to generate an innocuous looking file or USB stick with a trojan horse virus that once installed on your PC basically bypasses your firewall - by having your PC connect back to an attacker's machine. Once connected to the attacker, your PC can be used to "pivot" or scan your internal network for vulnerabilities.

SRTP, like TLS, is "hop by hop" and has to be established and maintained on each hop. Typically, the BCF provides a "media anchor", and thus the BCF terminates an incoming SRTP connection and originates another SRTP connection towards the PSAP, or another BCF. There may be one or more additional hops prior to the last hop to the call taker workstation.

Keying for SRTP can be provided by DTLS (the recommended way) or SDES (which many older systems have deployed). Where SDES is deployed, plans should be made to migrate to DTLS. Also consider the logging connections for media, which also need to be protected with TLS.

Persistent SRTP connections can be made in the same manner as persistent TLS connections to avoid crypto startup delay.

Consider how well your vendors can provide SRTP on video media. In some implementations, the cryptographic subsystems are unable to maintain secure video media connections. The standards require ALL media to be protected by TLS. It is possible to create systems that reuse keys between

hops such that the packets in media relays do not require decrypt (with one key) and re-encrypt (with another key) in order to maintain hop-by-hop security.

3.7 Data Rights Management

Once a user has been authenticated, authorization decisions must be made. NG9-1-1 standardizes the mechanisms by which access is granted to services or data using the XACML based Data Rights Management System described in NENA-STA-010. XACML is the OASIS standardized language for describing authorization policy. As described in NENA-STA-010, each XACML policy defines: a "target", which describes what the policy applies to (by referring to attributes of users, roles, operations, objects, dates, and more), and one or more "rules" to permit or deny access. Access is defined to mean some combination of:

- Read – the ability to retrieve a data object
- Update – the ability to modify an existing data object
- Create – the ability to create a new data object
- Delete – the ability to remove an existing data object
- Execute – the ability to execute one or more functions from a service.

Rules may "permit" or "deny" access.

XACML policies are stored in a policy store. The XACML "Policy Decision Point" can be inside the element or agency that has the "Policy Enforcement Point", or may be external to it.

Managing rights is rarely simple. Authorities must consider carefully how to provide appropriate levels of security while allowing agents to do their jobs. It is common to see roles inappropriately defined, or policies written too broadly to permit an uncommon occurrence to be handled simply, when it would be more appropriate to have more stringent security even if handling the uncommon event becomes more complex.

It is very difficult to manage security when applications or services insist on their own identities, authentication and/or access rights management instead of using the standard mechanisms. Customers should insist on the standardized mechanisms, and institute appropriate processes to create, monitor and audit access policies.

3.8 Dealing with Attack and Intrusion

A denial-of-service attack (DoS attack) or distributed denial-of-service attack (DDoS attack) is an attempt to make a machine or network resource unavailable to its intended users by overwhelming the resources of the target with more requests than it can handle. A specialized form of DoS is Telephony DoS (TDoS) which is an attack that creates many actual calls by corrupting many legitimate calling devices using some form of virus and causing it to place calls.

Avoiding DoS/TDoS is currently not possible. Such attacks should be considered "when", not "if". Network planners should utilize NENA 75-502, NG-SEC Audit Checklist. The Border Control Function is a critical network security element for IP multimedia services designed to effectively

manage sessions and protect core network elements from various types of distributed denial of service (DDoS) attacks including malicious and non-malicious signaling overload attacks. Depending on design, the BCF can be the demarcation point between trusted and untrusted network boundaries.

Border Control Function (BCF) as stated by NENA-STA-010: "sits between external networks and the ESInet and between the ESInet and agency networks. All traffic from external networks transits a BCF." The Border Control Function comprises several distinct elements pertaining to network edge control and SIP message handling. These include:

- Border Firewall
- Session Border Control

It is imperative that the border control function support the following security related techniques:

- Prevention
- Detection
- Reaction

Additionally, the entirety of the functional element may include aspects of the following:

- B2BUA (back-to-back user agent)
- Media anchoring
- Stateful Firewall

The BCF may be able to identify calls that may be part of a deliberate attack on the system. However, under normal conditions, we allow suspicious calls in, preferring to have a bad call show up to having a good call dropped. This means careful consideration must be taken with regards to the BCF configuration. There must be a refined manner in which suspicious calls are allowed, yet filter out what is confidently determined to be an attack. 9-1-1 authorities should be concerned with internal attacks and external attacks. Additionally, 9-1-1 authorities need to characterize overloads as an attack, whether malicious (intentional attack), or non-malicious (misconfiguration or malfunctions).

NG9-1-1 defines standardized functions that allow BCFs to mark calls as suspicious, which then allows the Policy Routing Function in the ESRP to route such calls differently than other calls. The standards also provide mechanisms for PSAPs to inform the BCF of so called "bad-actor" call sources that the BCF can filter. These mechanisms will be key to mitigating TDoS attacks and can also be used to filter malicious callers so long as a valid identifier of some form is available to identify them.

Of necessity, PSAPs will be connected, indirectly through the ESInet, to the Internet to accept calls. Each NG9-1-1 network should be designed under the assumption there can/will be Internet connectivity.

Generic DoS attacks from the service providers are not expected since the calls should be coming from highly trusted networks that are usually only able to carry telephony traffic. TDoS calls from

such network CAN be expected, because the networks do not have any way to differentiate. Since calls must be allowed from the Internet, the BCF cannot filter on the addresses of known call sources, but could arrange to have such calls differentially routed in the PRF. SBCs contain mechanisms to manage admission (number of calls allowed), but generally these mechanisms should not be used because they block legitimate calls when the source has TDoS traffic.

TDoS attacks may either be socially coordinated but are typically automated. In automated attacks another customer of the service provider (or one of their peers) generates calls using an IP-PBX or other device connected to both the Internet and the telephony network and corrupted by a virus. The calls may have a spoofed identity. If the BCF can identify a source of calls, even if that identity is spoofed, it may be able to filter calls from bad actors. There are some medium term activities aimed at preventing spoofing of ANI which will substantially improve the ability of NG9-1-1 systems to mitigate TDoS.

DoS attacks from the Internet are more complicated. Overprovisioning of bandwidth and server capacity will surely help you withstand smaller DDoS attacks. However, the size of DDoS attacks is growing, and it is not unusual to see DDoS attacks over 600Gbps now. DDoS attacks, even the largest seen DDoS attacks ARE being successfully mitigated. ESInets should be engineered to be able to mitigate the largest feasible attacks. However, attack size is an arms race, and it's possible that an attacker could muster more bad traffic than the mitigation capacity of the system. In such circumstances, some throttling of traffic may occur.

It is important to remember that a onetime drop of a SIP messages does not mean a call will be lost. SIP recovery mechanisms will re-initiate the call without user interaction. IP addresses that significantly exceed messaging rates and never continue the conversation they initiated (i.e. they don't respect session state) should be blacklisted. For example, a SIP INVITE flood tool will never respond to the "200 OK" responses sent by the BCF when the PSAP answers.

Floods coming from the Internet using a known signature of a scanning, reconnaissance, cracking, or flooding tool should always be blocked. There is no non-fraudulent use case for these. For example SIPVicious, a tool known for enumerating and cracking SIP credentials embeds the User-Agent descriptor "friendly-scanner" in the SIP messaging. No phone system or soft client does this, so it's guaranteed the request is fraudulent.

Recently, DDoS mitigation services have become available that provide an economical alternative to provisioning massive amounts of bandwidth to be able to mitigate the current large (600G+) sizes of attacks that are occurring. These services offload traffic to specialized hardware, coupled with the necessary bandwidth to mitigate these very large attacks. Authorities contemplating use of such services should ensure they are able to handle SIP as well as HTTP and DNS attacks. Activating these services when attacks occur involves cooperation with the ESInet IP providers.

Creation of a Public Safety Computer Emergency Response Team (CERT) is anticipated, and all BCF operators must arrange to receive alerts from the CERT and respond. It is essential that all BCF support organizations have trained staff available 24 x 7 x 365 to immediately respond to attacks and have the capability and training to be able to adjust the BCF to mitigate such attacks.

3.9 NAT Related Security Issues

Multimedia over IP communications run on the same IP networks as other IP applications. Almost all these IP networks deploy Firewalls and/or Network Address Translation (NAT) devices, for security. Unfortunately, the inherent characteristics of Multimedia over IP protocols are in conflict with most current mechanisms employed by Firewall and NATs, creating the need for NAT traversal solutions. The complexity of the problem together with the diversity of existing topologies means that different solutions are needed for different cases. NENA-STA-010 states "Network Address Translations (NATs) should not be used within an ESInet. Although NAT use within an ESInet is not recommended, NATs may be needed in specific deployments, and therefore all network elements must operate in the presence of NATs."

Network Address Translation (NAT) devices translate an IP address used within one network to a different IP address known within another network. One network is designated the inside network (for example, an enterprise LAN) and the other is the outside (for example, the Internet). Users on the inside network can see the outside network but the outside cannot see the inside users, as all communication with the outside network is via the translation device. Each outgoing or incoming request must go through a translation process, which is dynamic and transparent to the applications. This provides an opportunity to qualify the data by matching the source and destination of a packet in one direction to those of a packet in the opposite direction.

Typically, NAT devices let all the outgoing traffic traverse network boundaries. The addressing information of the packet is stored with a timeout. Packets flowing in an opposite direction that have matching addresses are allowed into the internal network. This technique of creating dynamic rules is called "pin holing". Typically, on outgoing packets a NAT device map local inside network addresses to one or more global outside IP addresses. On incoming packets the NAT device maps global IP addresses back into local IP addresses.

There are two flavors of Network Address Translation devices:

- A Network Address Translation (NAT) device allows an organization to use a range of private IP addresses when communicating within an inside network and to share a small pool of public IP addresses when communicating with an outside network.
- A Network Address Port Translator (NAPT) or Port Address Translator (PAT for short) device has a block of inside addresses and one or more outside addresses.

The nature of NATs and firewalls is such that connections from un-trusted or external devices are not allowed unless either a static policy, or "pinhole" is established to allow connections to a specific internal address or a connection has been previously initiated from an internal trusted address to external un-trusted entity.

In cases where a remote worker (call taker is not at the PSAP), is behind a NAT, it's publicly reachable IP address and port must be registered with a registrar service which maps a user's address of record to its location. The NAT problem for signaling can be solved by simply implementing a smart registrar/call agent function which does not save the contact address as presented by the device in the registration message but rather based on the real IP:port combination the message originates

from. Once registered, either the phone or the registrar function must maintain the communication channel open by sending keep-alive packets before the binding expires in the NAT device. The packets could be either signaling packets send by the device or IP packets sent by the registrar function. For example, with SIP it is possible for the Registrar function to instruct the endpoint to register more frequently.

Now, having a permanent communication path open between the registrar function and end user device, it is always possible to ring the device behind the NAT and to start negotiating a multimedia session. The only requirement, which is available in most of the endpoint devices, is to use symmetric signaling, that is the device must send and receive data on the same port number.

Each type of media stream may have one or more channels but each channel requires its own pinhole to be opened. This means that for the media stream to traverse the Firewall/NAT, the Firewall/NAT needs to open many UDP pinholes for each call session. Unfortunately, static provisioning of UDP pinholes leaves the network behind the Firewall/NAT exposed, defeating the primary mission of the NAT/Firewall. This is why NAT's present a significant security threat. A dynamic method that allocates pinholes based on authorized signaling connections is needed: this is a solution to the security threat. If, despite the recommendation, NATs are used within an ESInet, dynamic pin hole firewall techniques must be used to manage media connections. Signaling should be using TLS, and thus TCP, and NATs always dynamically open TCP pin holes.

The protocols used for voice and video call setup and signaling over IP are SIP based. These protocols use TCP as well as UDP for call setup and transport. TCP and UDP use port numbers to identify individual connections. Media transport addresses are embedded in the messages of the protocol (e.g., SDP). This results in an addressing conflict between the session control and transport layers. This commonly results in one-way audio, no audio, timeouts and drops. Older systems did not have mechanisms in the end devices to handle NATs in the path, thus Session Border Controllers devised methods for media traversal that incorporate a media relay function with an intelligent session agent that can discover the dynamically allocated source addresses assigned by the NAT device used for media and correct the end-to-end media transport address negotiation. Newer systems use a dynamic address/port negotiation mechanism known as "ICE", which is required to be supported by NG9-1-1 elements and the BCF will need to monitor the signaling and open pinholes for the media connections. Even with ICE, there may be reasons for the BCF to "anchor" media, which means it relays media. In that case the BCF does ICE negotiation and endpoints inside the ESInet see media addresses that land on the BCF.

3.10 Securing DNS

Domain Name Service (DNS) is one of the most critical, often overlooked, and prone to attack aspects of a network. There are a multitude of current hacks and vulnerabilities with more coming regularly. DNS is key to NG9-1-1 as it is fundamental to the processing and handling of any NG9-1-1 call.

In this discussion, a "meta" second level domain of "example.com" is used as the prototype name for 9-1-1 services. A particular regional ESInet subdomain is called "illinois.example.com" while a

particular PSAP might be “clintonco.illinois.example.com." Likewise, call handling function names should also be standardized, such as textincoming.clintonco.illinois.example.com

This document only includes DNS as pertaining to NG9-1-1 and does not address DNS for administrative functions (e-mail, web access, etc.).

3.10.1 DNS vs. Static IP addresses.

Devices connected to an IP network are always addressed by their IP address. IP addresses can be assigned by one of two mechanisms: a static IP address that is provisioned into the device or assigned dynamically by DHCP. To send a packet to any device, the sender must know the IP address. A convenient way to learn the IP address of a device is that it’s included in an IP packet. So if a device gets a packet from another device, it can always send a packet back to the sender, because it now knows its IP address. However, the sender would have to know the IP address of the device it needed to send the packet to before it can send one, and thus other mechanisms are needed to find IP addresses. Specifically, it can be provisioned with or it can discover a static IP address, or it can be provisioned with or discover a “domain name” and use the DNS to find the IP address of that device.

Use of static IP addresses is sometimes necessary. For example, the address of the DNS server pretty much has to be a static IP address (and the address is usually discovered by DHCP). Routers are almost always provisioned with static IP addresses. On the other end of the spectrum, a device like a call taker workstation is almost always dynamically assigned, and if needed, its address is discovered using DNS. While a device with a statically assigned IP address may have a DNS entry, a device with a dynamically assigned IP address that needs to be known by another device in advance (so it can send it packets before the other device replies) must have an entry in the DNS.

There are differences of opinion on how often static IP addresses should be used in preference to dynamically assigned IP addresses. It used to be the case that nearly every “server-like” device got a static IP address, and “user-like” devices got dynamically assigned addresses. That advice is largely obsolete, mostly because maintaining large numbers of statically assigned addresses is complex and error prone, and increasingly, “server-like” devices are getting dynamically assigned IP addresses. Within NG9-1-1, a device like the external ECRF, and the addresses (URIs) in it would need static IP addresses, but devices within the ESInet that are only reachable from within the ESInet could have dynamically assigned IP addresses and DNS entries.

Devices with dynamically assigned IP addresses as well as most devices with static IP addresses have an entry in the DNS. It is very common for there to be multiple “layers” of DNS, with a server acting as the authoritative server for interior resources but reaching out to an outer server for DNS names it cannot resolve. This allows, among other things, for such resources to be kept private. Certainly, a device that is assigned a private IP address, not accessible from the Internet should not have its DNS entry in the public DNS, if it needs to be reachable through the DNS, it would have its DNS entry in some local DNS server which does not publish such addresses to the Internet.

On the other hand, it is often recommended that devices that DO have a public IP address, but should not be generally accessible to the outside Internet not have a public DNS record. That is a form of

security by obscurity, and such a practice is generally not recommended. The problem with hiding the DNS entry is that trying to maintain the anonymity while allowing the authorized entities to get the entry is often complex and mistakes are very often made. Debugging can also be complicated by such attempts. If a device is addressable from the Internet, and it needs to be discoverable from outside the local domain, we recommend it be in the public DNS.

3.10.2 Authoritative Name Servers for externally addressable domains

Resources that are addressable from outside the ESInet must have the authoritative DNS servers available from the Internet. This can be accomplished in two ways: a DNS server can be put in the DMZ area behind the firewall portion of the BCF, or an external DNS service can provide the authoritative responses. DNS is regularly attacked (DDoS attacks specifically targeting DNS) because if you can stop legitimate users from getting the IP address of a service, you can prevent access to the service. Regardless of how authoritative DNS service is provided, it must be robust enough to withstand the largest feasible DDoS attacks, or the attacker will be able to stop legitimate 9-1-1 calls from getting to the ESInet. Third party mitigation services may be able to protect a server in the DMZ, but consider what must happen before an attack is discovered, and mitigated compared to using an external service that has the mitigation service itself.

Within the ESInet, authoritative DNS servers could be less robustly protected if they cannot be accessed from outside.

For DNS queries coming from outside the ESI Net, the DMZ ESI Net servers should be authoritative. For queries coming from within the ESI Net (including PSAPs), the internal ESI Net servers should be authoritative.

3.10.3 Exchange of DNS Information.

Internal DNS servers should be configured to transfer only the necessary DNS information to the respective external DNS server.

ESInet internal DNS servers should cache information from their respective external DNS servers as well as be authoritative for any devices within the ESI Net. In turn they should provide updates only for the minimum necessary devices to the external DNS servers, which should be authoritative for the “region.example.com” subdomain.

Making an assumption that there are multiple regional ESI nets and that they are interconnected through an internet, there should be an internet root level DNS function that is authoritative for the “example.com” TLD. This could be provided through various commercial vendors and should be both vendor and geographic diverse.

Outgoing DNS queries from the PSAP should be configured for request routing to query the region.example.com DNS servers for regional FQDN’s, the authoritative external DNS for any other .example.com FQDN’s, and their respective locale.

3.10.4 DNS Security

DNSSEC is required for any domain name in the ECRF and should be implemented elsewhere wherever possible. DNSSEC substantially improves the ability of a querier to trust the response they receive. DNS servers (except for the internal PSAP DNS which may be a service on a device such as an Active Directory server) should be implemented on dedicated hardware or virtual machines using hardened operating systems, with all unnecessary functions and protocols disabled. In addition, they must be protected by network security devices such as firewalls, intrusion detection and protection, etc. There is some value in disabling recursion in DNS servers, forcing the client to do iteration. DNS is one of the services that absolutely MUST be kept up to date with security patches.

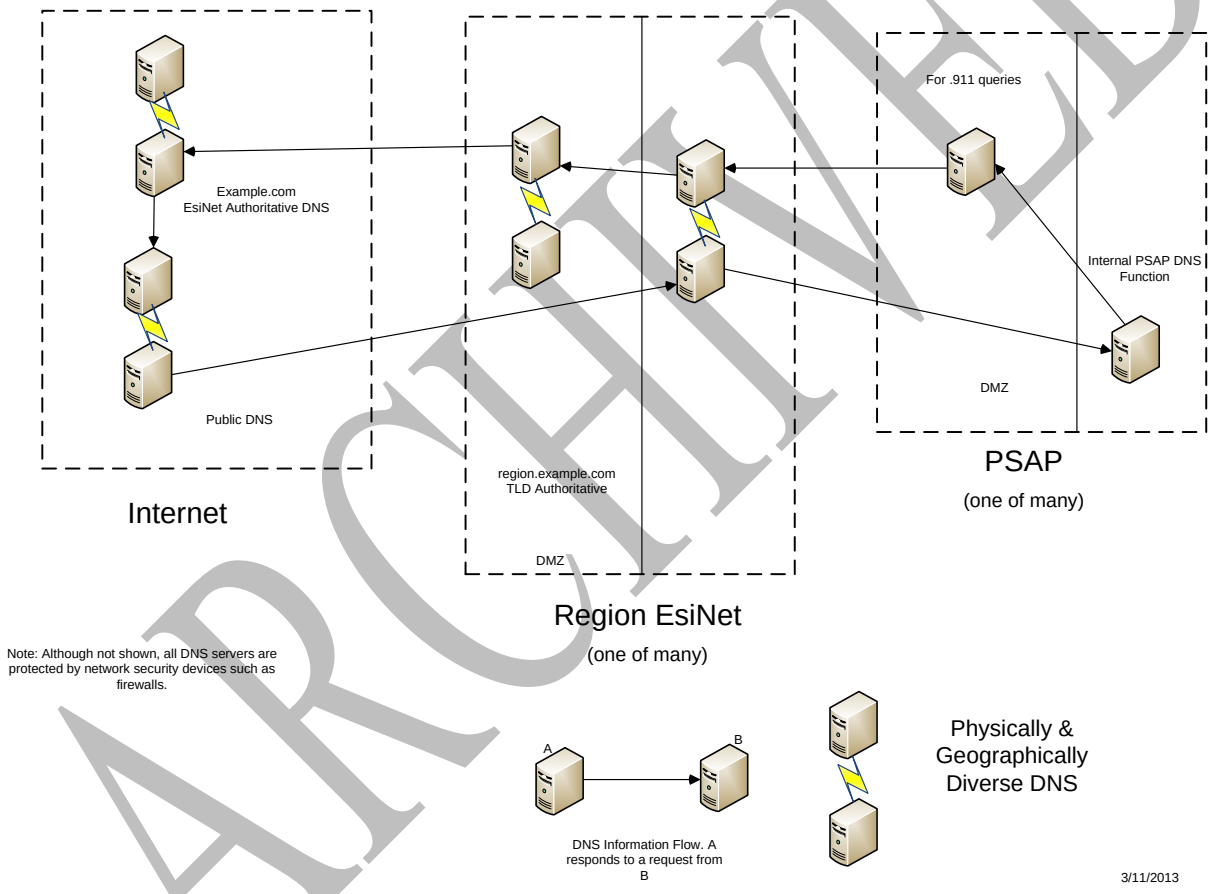

3.11 Security Issues in Connecting to Other ESInets

The necessity of interconnections between multiple ESInets is a core assumption in deploying the i3 architecture. Indeed, it is anticipated that every ESInet will be interconnected with at least one neighboring network. Thus, all ESInets will be interconnected at Layer 7. While these interconnections are essential to realize the goals of universal access and call processing free of constraining geographic boundaries, they also significantly impact the security of the interconnected networks. There is no trust assumed between interconnected ESInets. In addition to the issues

inherent in such interconnections, additional security impacts may stem from disparate operating authorities for each ESInet implementing differing, incompatible or non-compliant security policies in their respective ESInet.

Some of the areas of concern include:

1) Differing levels of trust assigned to individual networks as well as pieces of data transmitted through the network. Always use the security mechanisms that are roles based, as per NENA standards to enforce transitive vs. non-transitive trust relationships, and to convey information about the level and source of trust associated with a particular item or network. This is of particular importance when sharing data, as an agency may wish to share information with peers with which they have an explicit trust relationship with but not with the other agencies that the peers have a trust relationship with. In a properly implemented ESInet, there is no such thing as implicit trust: trust is controlled by an explicit policy, and rigidly enforced. For example, sending data by reference is a mechanism to allow the owner to maintain control access to its data.

2) Certain assumptions may be made about security arrangements when interconnecting with an i3 compliant network, but interconnection with a non-i3 compliant network may introduce additional security concerns and require more detailed security policy planning.

3) Some interconnection peers may require differing fundamental levels of security, for instance a regional civilian ESInet may need to interconnect with the ESInet providing 9-1-1 service for a military installation bordering the jurisdiction, but the two networks may have very different security requirements. Arrangements must be made in such cases to permit appropriate interagency call handling and collaboration while avoiding compromise of the more secure network.

4) Poor policy decisions, oversight, or technical implementation by administrators or technicians operating one ESInet may have significant negative effects on the security on the ESInets to which it interconnects. For instance, operators of an ESInet may make a decision to not deploy antivirus protection, or the technicians of that ESInet may fail to keep the deployed software updated. Because such lapses cannot be predicted in advance, no interconnection peer should be treated as fully trusted. The level of trust for an ESInet interconnection may be higher than the level of trust for interconnection to a traffic camera network, for instance, but every transaction should use standardized cryptographic based security mechanisms. Another area of significant concern is threat and attack management. This includes a number of separate issues, for instance:
    a. Mechanisms should be deployed for exchanging information about potential attacks in real time. What may appear to be an isolated event for one ESInet may emerge as a coordinated attack when events are correlated across several ESInets.
    b. Thresholds for declaring an attack must be coordinated across ESInets. Since the distinction between an attack and an unusual but actual event may difficult to

delineate, appropriate advance coordination will be essential to ensure that such events are handled appropriately.

c. Arrangements for mitigating an attack when detected, and in particular the authority to invoke such mitigating measures, must be made before the attack commences. Mitigation measures such as 'blackholing' a source IP address or even isolating a compromised ESInet have the potential to impact call handling capabilities for legitimate 9-1-1 calls as well as attack traffic, and therefore must be made appropriately yet swiftly. The responding appropriately to an ongoing event will be impossible without clear decision parameters and authority, and these arrangements will be especially critical in cases where an attack spans multiple interconnected ESInets.

The common thread in each of these areas is a need for prior coordination and agreement. If such decisions coordination, and resource management are only addressed when an attack is underway the response efforts will be significantly impeded, possibly critically so.

As recommended in NENA-STA-010, creation of a Public Safety Computer Emergency Response Team (CERT) is anticipated, and all BCF operators must arrange to receive alerts from the CERT and respond. It is essential that all BCF support organizations have trained staff available 24 x 7 x 365 to immediately respond to attacks and have the capability and training to be able to adjust the BCF to mitigate such attacks.

3.12 In SIP Trust Nobody

The ESInet is not a walled garden. No agency or element inside the ESInet should trust another agency just because it's inside the ESInet. The ESInet must be treated, for security purposes, as the open Internet. This does not imply that access to the ESInet is not controlled, it must be rigidly controlled. However, the trust model is based on cryptographic security, and not network access controls.

Ownership, common management, or trust relationships not based on the i3 cryptographic security mechanisms must not be used as a trust basis for i3 operations. Entities, agents and agencies have credentials, and these credentials are the basis for authentication and authorization for all i3 operations.

Of special note is the placement of border elements (BCFs) and gateways (LNG, LPG and LSRG). All calls, regardless of source must always pass through a BCF so that its call suspicion and filtering capabilities can be used. Denial of Service (TDoS and DDoS) attacks on PSTN elements, for example, is very feasible, and thus there are no circumstances where an LNG or LSRG can be placed inside an ESInet. They are always outside, and calls always pass through a BCF. Although only peripherally related to security, it should also be noted that LNGs may have to send calls to ESInets not near them when the local infrastructure is incapacitated (by attack or disaster), and thus assumptions based on local relationships may not always hold.

Also, BCFs should be placed between the ESInet and agencies like PSAPs and responders. Some consideration should be given to how these BCFs are provisioned, because there is as much concern

that a problem in an agency can corrupt ESInet operations as a problem in the ESInet can corrupt operations in the agency. Deployment of back-to-back BCFs is optional, based on needs and demarcation points. For example, the ESInet sends calls from its egress BCF to a PSAP BCF. Refer to NENA NTC Demarcation document for more details regarding physical and logical points of demarcation.

3.13 Security Issues in Connecting to the Internet

The Internet represents a vast pool of knowledge about nearly any topic imaginable, and much of that knowledge can be put to use by PSAP staff to improve response, make better decisions, and generally provide for improved service to the communities they serve. This notwithstanding, however, the Internet is also a fundamentally insecure network with many associated security risks.

NG911 systems will be engineered to accept calls from the Internet because not every service provider is capable of connecting via private interconnect to an ESInet. Therefore, appropriate security mechanisms must be implemented to handle such traffic.

Areas for particular attention include:

1) The necessity of providing access to the Internet via a facility that, in addition to more traditional firewall capabilities, inspects all traffic in detail, especially guarding against viruses, Trojans and other malware. Gone are the days where a specific file download was needed to infect a host with malware; in the absence of protective facilities just visiting a compromised website can cause a security compromise. Where protocols or applications unspecified by NENA-STA-010 are supported through the ESInet, then specific security controls or mechanisms must be implemented.
2) The same facility can be used to provide filtering based on site content. While giving the security benefit of rejecting traffic to and from many of the riskiest sites on the Internet, these capabilities also offer some protection against the legal ramifications of staff accessing such inappropriate content.
3) Patch management for systems with Internet access is especially critical. Software vendor updates and patches, especially for operating systems, should be deployed to all hosts as soon as deemed as safe and effective by the governing IT authority or, in the absence of such authority, when released by the software vendor.
4) Education, especially of end users, is a critical aspect of any security management plan. Areas for specific education include:
    a. The importance of maintaining secure credentials with strong passwords and of not sharing those credentials
    b. The importance of maintaining vigilance for unusual events – for instance, unexpected dialog boxes or error messages – and of promptly reporting such events to IT staff for investigation.
    c. The importance of every individual user having their own credentials and not using shared accounts or credentials.

3.14 Process and Audits

Achieving secure systems depends as much on processes as mechanisms. Processes need to be written, reviewed regularly and some form of record keeping is needed to assess whether a process has been followed. Staff needs to be trained on processes, and compliance to processes must be assessed. The best processes are those that are simple to understand and apply and do not get in the way of getting the job done. Security processes do not need to be cumbersome or onerous to be effective; indeed we often find that if the process is hard to use, creative alternatives are found and used that results in much worse security than a simpler process would achieve. The answer is NOT to hammer staff to obey the process. Change the process to achieve realistic goals without undue burden.

The very best mechanisms don't work if the processes to use them are shoddy or not uniformly applied. Auditing is an excellent mechanism to make sure processes are actually being applied correctly. Getting a qualified third party (which may be an auditor or just a knowledgeable person in another department) to audit the application of security practice will greatly improve the security of any facility or service and is highly recommended Every process should be audited regularly, with frequency determined by the level of concern, maturity of the process, and stability of the staff and mechanisms.

3.15 Other Protocols and Considerations

Properly securing an NG911 ESInet can be very daunting and time consuming. In addition to all of the above listed protocols and security measures, there are dozens of other protocols and hardware and software interconnections that need to be considered.

1. Protocol Versions -- Ensure that you are using the most secure up to date version of the protocol.
    a. IPv4 versus IPv6
    b. SNMP v1, v2, v2c, v3...
2. Replacing protocols with more secure protocols
    a. Replacing Telnet with SSH
    b. Replacing FTP with SFTP
3. Remove unnecessary and outdated protocols
    a. Uncheck IPv4 from servers that only need to communicate to IPv6 networks
    b. Upgrade all SNMP v1, v2, only use v2c when necessary, preferred v3
4. Use Firewalls to filter between networks
    a. Filter unneeded data/protocols from the WAN. e.g, Telnet
    b. Consider where the data is coming from, where it is going to, and what value it has.
5. Consider the ramifications of filtering protocols
    a. IPv6 requires ICMP for packet fragment detection and correction.
    b. Troubleshooting (both locally, and across the ESInet)
6. Keep hardware and software up to date

a. Hardware maintenance
b. Software maintenance
c. End of Life -- If the device is no longer supported by the manufacturer, it should be replaced

7. Proper Security Audits
   a. Limit physical access
   b. Disable unused ports on switches and routers
   c. Annual review of accounts
   d. Proper password expiration
   e. Review the FBI CJIS Security Policy for additional information

4 Recommended Reading and References

- NENA 75-001, Security for Next Generation 9-1-1 Standard
- NENA 75-502, Next Generation 9-1-1 Security Audit Checklist Information Document
- NENA-STA-010, Detailed Functional and Interface Standards for the NENA i3 Solution

5 Previous Acknowledgments

Not applicable – initial document.

Appendix A: Security Checklist Table

	Security Element								
Functional Element	TLS	SRTP	Single Sign on	Credentials traceable to PCA	Patches	Communicate with policy mechanisms	Call path	Databases	External (to ESInet) interface
LNG			Note 1	Note 1				n/a	
BCF								n/a	n/a
ESRP		n/a						n/a	n/a
ECRF		n/a					n/a		
LVF		n/a					n/a		
LPG								n/a	
Call Handling								n/a	n/a
Dispatch		n/a					n/a	n/a	n/a
IDE		n/a					n/a	n/a	
Outgoing Alert		n/a					n/a	n/a	n/a
Incident Handling		n/a					n/a	n/a	n/a
IMR								n/a	n/a
Map database		n/a					n/a		n/a
MIS		n/a					n/a		n/a
RMS		n/a					n/a		n/a
Mobile data		n/a					n/a	n/a	n/a
Logging Service		n/a					n/a		n/a
Media recording interface									n/a
Radio interface								n/a	n/a
Bridge								n/a	n/a
Agency locator service		n/a					n/a		n/a
ADR		n/a	n/a	n/a			n/a		
Note 1: If LNG is provided by 9-1-1 authority then SSO and credentials traceable to PCA apply.									
Note 2: Nearly all FE's will need to enforce data rights management									