Exploiting Lawful Intercept to Wiretap the Internet

Tom Cross
IBM X-Force
tcross@us.ibm.com

Abstract

Many network equipment manufacturers have incorporated interfaces into Internet routers and switches that are designed to facilitate legally authorized wiretapping by law enforcement. If these interfaces are poorly designed, implemented, or managed they can provide a backdoor for attackers to perform surveillance without lawful authorization. Most lawful intercept technology is proprietary and difficult to peer review. Fortunately, Cisco has published the core architecture of it's lawful intercept technology in an Internet Draft and a number of public configuration guides.

This paper will review Cisco's architecture for lawful intercept from a security perspective. We explain how a number of different weaknesses in its design coupled with publicly disclosed security vulnerabilities could enable a malicious person to access the interface and spy on communications without leaving a trace. We then provide a set of recommendations for the redesign of the interface as well as SNMP authentication in general to better mitigate the security risks.

1. Introduction

Operators of large public internet networks inevitably face requests from law enforcement to access communications carried on their networks. Special network interfaces are often built to facilitate that access. In some contexts network operators are required by law to create convenient interfaces, but these interfaces are often deployed even when they aren’t required because they can simplify the process of facilitating law enforcement access and make it minimally disruptive to network operations.

These interfaces are controversial. If they are not well-protected there is a risk that they could be hijacked by third parties and used to perform surveillance without authorization. [1] Because of this risk, the security of lawful intercept systems is of obvious public interest. The IETF has published a policy on wiretapping which simultaneously argues that the IETF is not an appropriate standards body for developing lawful intercept interfaces, for several reasons including the security risks those interfaces potentially pose, but the IETF nevertheless encourages the open publication of lawful intercept architectures to facilitate peer review. [2]

In keeping with this approach, Cisco has published the core architecture of its lawful intercept technology in an Internet Draft and a number of public configuration guides. This is good for two reasons. First, it enables the general public to see and understand how wiretapping is performed with Cisco routers. Second, it allows the security community to peer-review their approach to protecting this interface from attack.

That peer-review is the basic purpose of this paper. We review Cisco's architecture for lawful intercept and explain the approach a bad guy would take to getting access without authorization. We identify several aspects of the design and implementation of the Lawful Intercept (LI) and Simple Network Management Protocol Version 3 (SNMPv3) protocols that can be exploited to gain access to the interface. In particular, an implementation flaw in SNMPv3 is discussed in the context of LI which could have provided direct access to the LI interface to malicious users without a password until it was patched. Finally, we provide recommendations for mitigating those vulnerabilities in design, implementation, and deployment.

Ultimately, our intent is to raise questions about the prevailing approach to protocol design, wherein protocol standards provide maximum flexibility to their implementers, who in turn provide maximum flexibility to the operators who deploy the end system. This is the natural result of the economic forces involved in the design and deployment of network technology, but some deliberate effort may be required to counter it. The negative consequence of this approach is that operators are left with an array of different choices for how to deploy a system, some of which are bad. Inevitably, some operators will make those bad choices. If Internet protocols were designed and implemented to steer operators toward better deployment choices, security would benefit.

2. The Cisco Architecture for Lawful Intercept in IP Networks

Cisco is one of the only companies to have published technical details about their solution for implementing lawful telecommunications intercept. The core of this solution is described in IETF RFC 3924. [3] In spite of the IETF's rejection of Internet Standards for wiretapping, the RFC Editor decided to publish this document at their discretion. It begins with a clear note from the Internet Engineering Steering Group that this RFC is not a candidate for any level of Internet Standard and the IETF makes no representation as to its fitness for any purpose.

This document describes a general reference model for telecommunications intercept which includes a number of architectural components and the interfaces between those components. Figure 1 provides an overview of this architecture. When the Law Enforcement Agency (LEA) wishes to perform surveillance they contact the Service Provider's Lawful Intercept Administration, an organization that verifies the LEA's legal authorization to access the content they seek. The Lawful Intercept Administration uses the Mediation Device (MD) to provision surveillance. The Mediation Device crafts an Interception Request based on the content to be collected and sends it to the Intercept Access Point (IAP), which is typically a network device like a router or a switch which has access to network traffic. In response to this Interception Request the IAP collects the network traffic that the Law Enforcement Agency is interested in, and sends it back to the Mediation Device. The MD reformats that information (if required) and retransmits it on to the LEA.

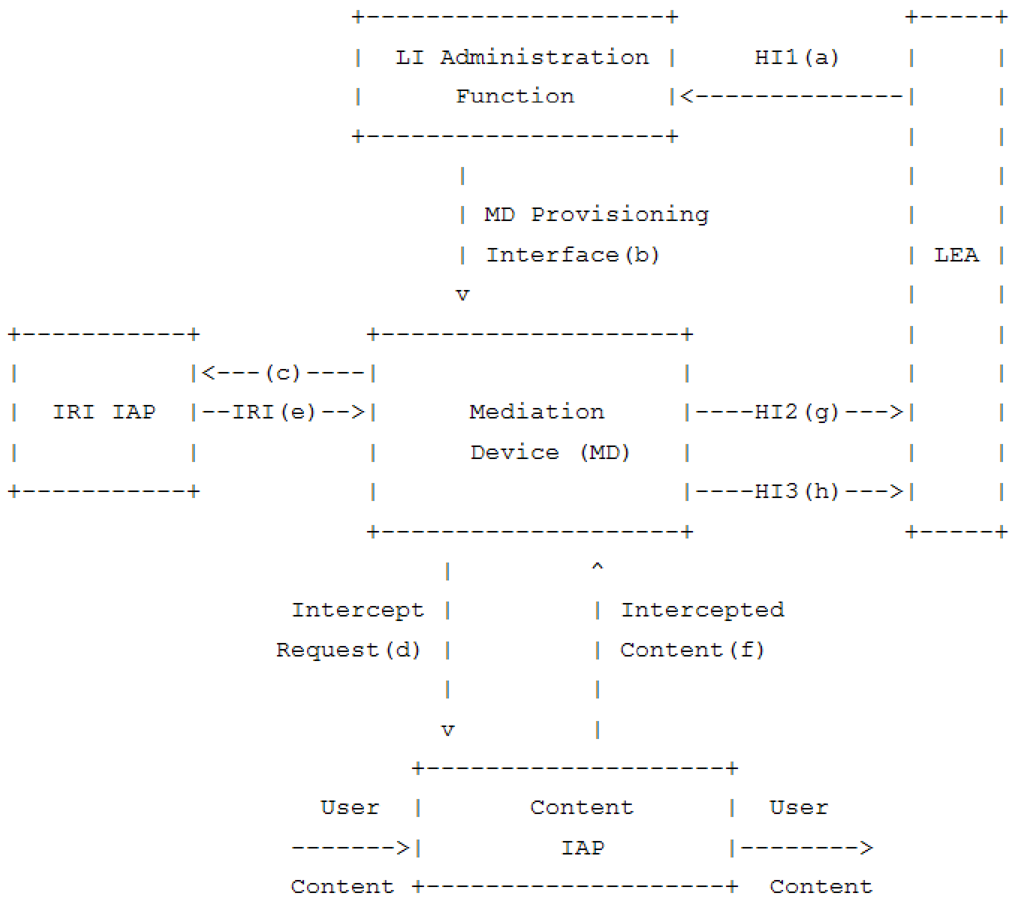

```
                +--------------------+         +-----+
                |  LI Administration |   HI1(a)  |     |
                |      Function      |<------------|     |
                +--------------------+           |     |
                          |                      |     |
                          | MD Provisioning      |     |
                          | Interface(b)         | LEA |
                          v                      |     |
+-----------+     +--------------------+         |     |
|           |<---(c)----|              |         |     |
|  IRI IAP  |--IRI(e)-->|   Mediation  |----HI2(g)--->|     |
|           |     |     Device  (MD)   |         |     |
+-----------+     |                    |----HI3(h)--->|     |
                  +--------------------+         +-----+
                        |         ^
              Intercept |         | Intercepted
             Request(d) |         | Content(f)
                        |         |
                        v         |
                   +--------------------+
              User |      Content       |  User
           ------->|        IAP         |-------->
           Content +--------------------+  Content
```

Figure 1: Intercept Architecture

Source: RFC 3924
Copyright (C) The Internet Society (2004).

The document goes on to explain that the Interception Request is implemented using Simple Network Management Protocol Version 3 (SNMPv3). SNMP is an Internet standard protocol which is used for monitoring and managing devices on a computer network, such as routers, switches and servers. In most deployments, the protocol involves sending UDP messages, each of which may consist of a single packet, between network devices and a monitoring server. SNMPv3 is the latest version of SNMP. One of the most important advantages of SNMPv3 over prior versions of the protocol was supposed to be improved security, including built-in encryption as well as a vastly improved authentication system that allows devices to verify the authenticity of messages without sending passwords over the Internet in the clear.

The Interception Request is a SNMPv3 message which accesses a TAP-MIB. A MIB or Management Information Base is a collection of objects which describe a capability that can be managed over SNMP. Cisco's specific TAP-MIBs are described in separate documents. [4] The TAP-MIBs allow the sender of the Interception Request to specify many parameters for the interception, such as the source and destination addresses, port numbers and protocols of the traffic to be collected, whether or not notification should be enabled, and the IP address and port number of the Mediation Device that the collected traffic should be forwarded to. The TAB-MIBs also specify how to format the collected data. The options include a proprietary UDP packet format from PacketCable, a Real-Time Protocol stream, a TCP session, or an SCTP session. Basically, all of the details of the interception are controlled by the contents of the SNMPv3 Interception Request message.

Like all modern IETF documents, RFC 3924 contains a security considerations section. This section argues that message integrity checking, a feature of SNMPv3, is necessary for this application. The SNMPv3 View Based Access Control Model is used to ensure that the TAP-MIB is only accessible to particular SNMP users who have been granted access specifically for the purpose of performing intercepts. The document also states that "Privacy and confidentiality considerations, may also require the use of encryption." However, no specific encryption methodology is specified. The Internet Draft describing Cisco's TAP-MIB goes further, specifically recommending that IPsec ESP be employed to encrypt traffic. [9] (IPsec ESP is a standard for encrypting Internet traffic.)

3. Attacking Lawful Intercept

If an unauthorized person crafted an SNMPv3 Interception Request and successfully sent it to a router or switch supporting this interface, that person could wiretap communications. There are a few of pieces of information that the attacker would need to collect in order to craft an Interception Request that would be accepted. The attacker would have to obtain the correct SNMPv3 EngineID, EngineBoots, and EngineTime values for the device. Those values are used by SNMPv3 to prevent requests from being replayed. They can be obtained directly from the device via an unauthenticated SNMPv3 request. The attacker would also have to correctly guess or determine the SNMPv3 username and password that has been granted access to the TAB-MIB. However, an implementation flaw in SNMPv3 could be exploited to access the interface without the correct password. Once successful, the attacker could choose to capture any traffic on the device and route that traffic to any destination IP and port on the global Internet, over UDP, TCP, or SCTP.

Although there are features of these protocols and the network devices that they run on which are designed to thwart an attack of this sort, there are also several aspects of the design and implementation that work in the attacker's favor. In particular, the six aspects discussed below leave us with the impression that this interface could be successfully attacked in many real world deployments.

3.1 The Susceptibility of SNMPv3 to Brute Force Credential Discovery

In order to send an authenticated SNMPv3 request that will be accepted by the device that receives it, the sender must know an appropriate username and password. Unfortunately, the User-based Security Model for SNMPv3, described in RFC 3414, [5] makes it easy for an attacker to guess valid credentials by brute force. When an authentication failure occurs, the standard requires the receiving device to send a Report back to the sender. The Report includes different information when the authentication failure was due to a bad username (usmStatsUnknownUserNames) versus a bad password (usmStatsWrongDigests). Therefore, the attacker can try different usernames until the

Report received from the device indicates a bad password, and then try different passwords until the message is accepted. It is not possible to fix this weakness and comply with RFC 3414.

3.2 The Password Implementation Vulnerability in SNMPv3

Attackers may not need to guess a password via brute force due to an implementation flaw in SNMPv3. [6] This vulnerability allowed an attacker to access any password protected SNMPv3 service running on an unpatched device, without knowing the password, by sending 256 messages, one of which was guaranteed to be accepted as valid.

To understand how this attack works it helps to have a little background on SNMPv3 authentication. Two authentication protocols are currently defined for SNMPv3, HMAC-MD5-96 and HMAC-SHA-96. [5] Both protocols operate in a similar fashion, but use different hashing algorithms. In both cases a secret password is first established and shared between client and server. The client combines this password with the unique EngineID value of the server's SNMPv3 service to create a secret key. When the client wants to send a message, it calculates a cryptographic hash using that secret key and the contents of the message. According to RFC 3414, each message includes the first 12 bytes of its cryptographic hash. When a server receives a message, it calculates the same cryptographic hash and then compares the first 12 bytes with what was received. If both are the same, the client must know the correct password, and so the message is considered authentic.

Obviously, the server must perform a comparison operation in order to determine whether the cryptographic hash that it calculated is the same as the cryptographic hash that it received. Often in C programs such a comparison is implemented with a function like memcmp, which takes pointers to two strings, and a length in bytes to compare. Of course, the client only transmits the first 12 bytes of the cryptographic hash, and the hash calculated by the server may be longer, so many SNMPv3 implementations used the length of the cryptographic hash received from the client in the memcmp call. [7]

The standard requires that the server throw out any message containing a hash shorter than 12 bytes, but many implementations did not do this. The consequence is that an attacker need only send a single byte as the message digest, and if that byte is the same as the first byte of the digest calculated by the server, the message is considered authenticated. Attackers could therefore access password protected SNMPv3 services by sending 256 messages, each providing a different guess for the first byte of a valid message digest. One of these 256 messages would be accepted as valid and processed.

There are many potential attacks against network devices that have this vulnerability, depending on what they are configured to allow SNMPv3 users to do. For example, complete configuration files can be loaded on to Cisco routers over SNMP if the correct MIBs are enabled. [9] However, the case of lawful intercept is particularly interesting because network operators may be compelled to enable it, and may not able to balance the benefits of these features against the security risks they pose.

This vulnerability was publicly acknowledged in June of 2008 and assigned the ID number CVE-2008-0960 by Common Vulnerabilities and Exposures. [6] However, the problem had existed in some SNMPv3 code bases for over 6 years prior to that date.

[8] Numerous vendors were acknowledged to be vulnerable and shipped product updates to fix the problem, including various Linux distributions, Solaris, and Mac-OSX, as well as network devices from Juniper and Cisco. [6]

Fortunately, most versions of the Cisco IOS software that support the Lawful Intercept Interface are not vulnerable to the SNMPv3 password vulnerability according to Cisco's Security Advisory. [10] However, there are a few exceptions. For example, the Lawful Intercept Interface for Cisco 10000 series routers is supported by IOS version train 12.3(7)XI. [11] According to Cisco's Security Advisory, the first non-vulnerable version of IOS 12.3(7)XI is 12.3(7)XI8a. The "Open Caveats" for 12.3(7)XI2 listed on Cisco's website mentions known problems with the Lawful Intercept Interface, confirming its presence in a vulnerable version of the software. [12]

Cisco describes their 10000 series routers as "ideal for carriers deploying IP/MPLS services to broadband and private-line customers" and "the industry-leading edge router for service providers who require subscriber awareness for triple-play, broadband, and wholesale access." [13] Certainly, routers like that have access to the personal communications of many thousands of consumers. Furthermore, they may be used to provision point to point Virtual Private Networks that interconnect enterprise and small business corporate offices. [14] These so called "IP-VPNs" carry sensitive, internal network traffic that is logically positioned behind corporate firewalls, and is often sent without encryption as service provider networks are generally thought to be invulnerable to unauthorized surveillance.

3.3 The Lack of Audit Trails

Our attacker can not only access the Lawful Intercept interface without authorization, he may be able to do it without leaving a trace behind him. SNMP traps (or informs) might have enabled network administrators to detect misuse of this interface by reporting authentication failures and unauthorized access to a network monitoring facility. Cisco's Configuration Guide for Lawful Intercept advises administrators to configure intercept devices to send SNMP trap notifications to the Mediation Device when authentication failures occur. [11] Cisco's documentation implies that these traps will be sent "for packets with an incorrect SHA/MD5 authentication key or for a packet that is outside the authoritative SNMP engine's window (for example, outside configured access lists or time ranges)." [15]

Unfortunately, this does not appear to be the case. None of the IOS versions we tested sent authentication failure traps for SNMPv3 messages with the wrong username, password, or Engine values. Authentication failure traps were only generated for SNMPv3 requests if they came from a source IP address that was blocked by a group access list. We reported this issue to Cisco assuming that it was an implementation flaw, but Cisco concluded that this behavior was expected and that it was the documentation which was, in fact, incorrect.

Cisco routers can also generate SNMP traps that notify the Mediation Device of the current status of an Interception Request. That feature might have provided a powerful audit trail in the event of a compromise. If a Mediation Device received a status notification for an Interception Request that it did not initiate, that would be an immediate cause for alarm. Unfortunately, Cisco allows these notifications to be disabled through the very SNMP TAP-MIBs that are used to request an intercept. [11] Therefore, it is trivial for attackers to disable these notifications in the course of setting up their unauthorized surveillance.

3.4 The Flexibility of the Output Stream

The person crafting an Interception Request can specify any destination host and port number to send the collected traffic to, and can choose from several different formats for that data, including a proprietary UDP packet format from PacketCable, a Real-Time Protocol stream, a TCP session, or an SCTP session. This great flexibility means that once an attacker has successfully issued an Interception Request, he or she can receive the collected traffic from anywhere on the Internet. For example, the attacker could specify that the traffic be sent as a Real-Time Protocol stream on UDP port 53. To most packet filters and firewalls this traffic stream would be indistinguishable from Domain Name System (DNS) traffic, which is usually not filtered on Internet networks.

3.4 The Susceptibility of the Interface to Packet Spoofing

Most Cisco devices allow network administrators to block traffic on a network by indicating the source and destination IP addresses and port numbers for offensive traffic in an access list. Access lists are often the first tool network administrators turn to when dealing with a security problem.

Unfortunately, access lists have only limited effectiveness for protecting Lawful Intercept because Interception Requests can be easily spoofed. An Interception Request can be sent in a single UDP packet, and it can specify any destination host and port number to send the collected traffic to, and so theoretically the attacker does not need to be able to receive responses from the targeted device directed back to the IP address that the Interception Request came from.

However, the attacker’s job is somewhat complicated by the need to collect the EngineID, EngineBoots, and EngineTime values for the target device. As previously mentioned, those values can be obtained through an unauthenticated SNMPv3 request to the target device, but that request cannot be spoofed because the attacker would need to be able to see the response.

In theory it might be possible to collect these values without directly requesting them. On Cisco routers the EngineID defaults to a value which is generated based on the MAC address of the first interface on the device, a value which might be obtained from the network in certain circumstances. [16] EngineBoots and EngineTime both start counting from zero, and so they exhibit a bias toward low numeric values which might make them relatively easy to guess. However, attacks are far easier to perform if these pieces of information can be obtained from the targeted device through a normal SNMPv3 transaction that is not blocked by an access list.

The Cisco Security Advisory for the aforementioned SNMPv3 password vulnerability recommends using access lists to block all SNMP packets destined to a vulnerable router from untrusted IP addresses. [10] This work around is a good suggestion, as it blocks access to the Engine values from outside of the provider’s network, but it still leaves a great deal of exposure. In a real IP network SNMP is used for many purposes and a number of different systems on different networks may need to communicate with a particular router or switch via SNMP. The IP ranges for all of these devices would need to be allowed. These trusted networks may be subject to their own computer security vulnerabilities, and they may be legitimately accessible to service provider employees who are not authorized to perform intercepts.

In addition to defining an access list that limits the general SNMP service to trusted source addresses, Cisco provides a way to limit access to the lawful intercept feature specifically. In the course of configuring the Lawful Intercept Interface, the administrator has to define an SNMPv3 user group with access to the TAP-MIBs. An IP access list can be applied to that user group, limiting lawful intercept access to requests coming from the actual IP addresses of the Mediation Device. [17]

This is not a perfect mitigation either. An attacker with access to one of the regular SNMP monitoring networks could obtain the EngineID, EngineBoots, and EngineTime values with a regular transaction, and then spoof an Interception Request with those values from the Mediation Device's address. But, the attacker would have to know the correct address of the Mediation Device, and learning it might present a barrier to some attackers. Unfortunately, the ability to specify a user group access list was not mentioned in the Cisco Security Advisory for this issue nor in the SNMP hardening guide that advisory referenced. [10]

There are a number of anti-spoofing features and strategies which are described in Cisco's Applied Mitigation Bulletin [18] for the SNMPv3 password vulnerability. In general, the intent is to block packets that have an impossible source address given the direction they are coming from in a network. In order for this approach to be effective it has to be used consistently within a network. Internet security would be in much better shape were that the case, but in practice filters like these are only used in the most well run IP networks.

3.6 The Lack of a Requirement for Encryption

Encryption ought to be an important part of a lawful intercept capability. As the Cisco Lawful Intercept RFC (RFC 3924) points out in its security consideration section, there is a need to protect both the contents of the information that is being collected as well as the specific identities of the intercept subjects. [3] If Interception Requests and their corresponding traffic are being sent across networks in the clear, there is a risk that this information could be exposed. Furthermore, proper use of encryption can help protect the interface against unauthorized use.

However, Cisco does not require that encryption be configured for lawful intercept nor do they recommend a specific approach. Cisco's Intercept Architecture Guide states that "Because of privacy and confidentiality considerations, the architecture should allow for the use of encryption. Although encryption is not necessarily a requirement, it is highly recommended and may be a requirement in some LI deployments." [19] As some approaches to encryption work better than others, it makes sense to consider them carefully.

3.6.1 SNMP Encryption

The Cisco Security Advisory for the SNMPv3 password vulnerability recommends enabling SNMPv3 encryption as “a short-term workaround for users who are unable to upgrade in a timely fashion.” This mitigation effectively prevents the password attack, however SNMPv3 encryption is not required for lawful intercept. Cisco’s Configuration Guide [11] for Lawful Intercept on 10000 Series Routers states that “users must have authPriv or authNoPriv access rights to access the Lawful Intercept MIBs.” AuthNoPriv means without privacy or without encryption. Also, all of the configuration examples provided in this guide show authNoPriv access rights being configured.

Even when SNMPv3 encryption is enabled, weaknesses in the Lawful Intercept interface can facilitate attack by malicious insiders who know the correct authentication credentials. The threat of malicious insiders is significant because the interface is so susceptible to packet spoofing, and because it is possible to disable SNMP trap notifications. An insider with access to the correct SNMP encryption and authentication passwords could use the interface to monitor communications from anywhere on the Internet, in spite of any access lists that might be in place, and unbeknownst to the legitimate operators of the network. A secondary layer of security is really needed to prevent this kind of misuse.

To be fair, Cisco routers and switches do a good job of protecting the secrecy of SNMPv3 passwords. The passwords are discarded once authentication and encryption keys are derived from them, and the keys are never included in router configuration files. Furthermore, Cisco seems to have been concerned that the keys for lawful intercept might be recovered from NVRAM, and so in some models those keys are only stored in volatile memory and their passwords must be reentered manually every time the device is restarted. [20] However, passwords can be misused regardless of how well protected they are, and so there remains a need for secondary access control and audit capabilities.

3.6.2 IPsec

As previously mentioned, the Internet Draft for Cisco’s TAB-MIB suggests that IPsec ESP be used to encrypt these transactions. This is good advice, as the Lawful Intercept Interface offers no other facility for encrypting the collected traffic on its way back from the router to the Mediation Device. Law Enforcement agencies using this system may insist that this traffic be encrypted in order to reduce the risk that their targets will become aware of the fact that they are under surveillance.

On its own, a simple IPsec Security Association between the Mediation Device and the IAP also won’t mitigate an attack on SNMPv3 authentication. The router will accept unencrypted requests from other source addresses, and when the attacker directs collected traffic back to his or her address, that traffic will also be sent in the clear. However, IPsec coupled with proper access-lists can be effective. Attackers cannot easily spoof IPsec encrypted requests because the session keys used to encrypt them are relatively difficult to obtain, even for malicious insiders who know the shared secrets used to establish those sessions.

The general infrastructure access lists recommended in the Cisco Security Advisory for the SNMPv3 password vulnerability will not work, because they focus on closing off general SNMP access from untrusted source addresses. As previously mentioned, there are usually multiple SNMP monitoring hosts which require access to a network device. In practice these SNMP monitoring networks are unlikely to have IPsec security associations with the devices they are monitoring. Unencrypted Interception Requests could be sent or spoofed from those source addresses.

The best way to effectively limit abuse of this Lawful Intercept Interface is to couple an IPsec Security Association with a group access list that limits TAB-MIB access to requests coming from the IP address of the Mediation Device. This configuration prevents outsiders from attacking SNMPv3 authentication by requiring that requests be encrypted and it prevents malicious insiders who know the correct passwords from spoofing Interception Requests from networks they control. As notifications can be disabled there are still avenues for abuse in this configuration, but it is as secure as it can be given the design of this architecture.

4. Improving the Lawful Intercept Interface

Before we set upon the task of addressing some of the weaknesses raised in this paper we must first consider where the weaknesses should be addressed and who should address them. Many of these weaknesses stem from a desire on the part of protocol designers and implementers to serve their customers by providing them with the greatest possible flexibility. For example, the interface can be used regardless of whether or not encryption is configured, notification can be turned on and off by the user, and the Mediation Device can be changed without the need to reconfigure a router.

It is the natural economic interest of tool makers to empower the customers they serve as much as possible. But sometimes all of that flexibility can increase the likelihood that a protocol will be deployed in a way that is insecure. Modern IETF documents include a Security Considerations section in which protocol architects are asked to consider various attacks against their design and how implementers and operators can mitigate them. Unfortunately, if a particular attack can be mitigated through proper configuration, this is often viewed as sufficient. Protocol designers are loath to consider changes to protocols to address a security vulnerability that can be mitigated through proper configuration and management.

The Session Initiation Protocol (SIP) provides a simple example of the dangers involved in pushing responsibility for security away from protocol design and into implementation and configuration. SIP is an IETF standard protocol used to implement Voice over IP. SIP can be used to initiate telephone calls with a single, unauthenticated UDP packet, which can be easily spoofed. Attackers can leverage SIP devices on the Internet to perform traffic amplification denial of service attacks, by turning one spoofed UDP packet into a large stream of encoded audio data headed for their victim. Using this technique it's easy for attackers to saturate their victim's network connectivity.

The Security Considerations section of the SIP RFC addresses this problem by stating that "[User Agents] and proxy servers SHOULD challenge questionable requests" by forcing them to authenticate. [21] Authentication can help prevent spoofing in SIP by requiring an interactive exchange of packets. However, many SIP operators want to allow people to use their service without first establishing a username and password. The RFC offers to resolve this problem by suggesting that operators create an account with the username "anonymous" and a blank password that clients are forced to authenticate to, instead of allowing unauthenticated single packet requests. In practice, operators rarely do this and implementations usually don't point them in the right direction.

One reason that SIP operators aren't clamoring to enable authentication for anonymous SIP is that they don't directly bear the cost associated with traffic amplification attacks. The victims of these attacks are a third party. The SIP service is merely a conduit through which the attack takes place and many operators are not aware that their service can be abused in this way. In economics, this situation is referred to as a negative externality.

A negative externality is a cost that a third party must bare as a consequence of a voluntary and mutually beneficial economic transaction between two other parties. Negative externalities in the real economy can be difficult to resolve without regulation. The most obvious example is pollution, which can impact people who do not directly benefit from the industrial production that causes it. Our society manages pollution by creating rules that require producers to take steps to limit it even though those steps are not in their immediate best interest.

In the context of a network protocol, one way to manage a negative externality is to limit the flexibility that can give rise to it, in the protocol design and implementation. If the SIP protocol design, at a low level, required an interactive process even for anonymous messages, packet spoofing would be more difficult. SIP implementers could also reduce the risk by designing their software so that it is difficult to configure in such a way that unauthenticated, single packet requests are honored. Either of these steps would help reduce the likelihood of insecure SIP deployments on the Internet.

A negative externality can also exist in Lawful Intercept protocols. These protocols are designed to mediate a relationship between law enforcement and a network operator. The users of the network are a third party. Design and deployment decisions that facilitate the needs of the network operator or law enforcement may have a negative impact on those network users by exposing their traffic to attackers.

Although our analysis of potential attacks on Cisco's Architecture for Lawful Intercept identifies configurations that can prevent the attacks, we are not confident that every deployment of this protocol on the public Internet has used those configurations. The best mitigation involves coupling IPsec encryption with an unusual user group IP access list that is not described in the configuration documentation for Lawful Intercept, nor in the security advisory for the SNMPv3 password vulnerability, nor in the SNMP hardening guide that security advisory referenced. [10], [11] Although IPsec encryption is mentioned in the Security Considerations section of the Internet Draft for Cisco's TAP-MIB, advice about specific encryption configurations did not make it into Cisco's documentation and administrators are free to deploy the interface with no encryption at all. You can draw an observation from this that can be applied to everyone working with Internet protocols: It is the responsibility of each person involved in the process of moving from design to implementation to operation to ensure that the right information about secure implementation, configuration, and use is carried forward and does not remain unnoticed at the bottom of an expired Internet Draft.

Nevertheless, it is our view that good configuration guides are not enough. Changes ought to be made to these protocols that reduce the risk that they can pose to user's security when they are poorly configured, because we cannot assume that every network operator uses all the best security practices. As Sun Microsystems Fellow Radia

Perlman once said “An issue that separates real network protocol design from, say, a theoretical algorithm, is realizing that the components don't always behave as they should. A network has to be designed so that components that are configured incorrectly, or misbehaving in various ways can't do too much damage.” [22]

We present a set of recommendations for changes to the User-based Security Model for SNMPv3, the Lawful Intercept protocol, as well as advice for network operators deploying this system as it is currently designed. We recognize that there are almost certainly use cases and application considerations that we are not aware of that may make adoption of some of these recommendations impractical. However, adoption of at least some of these recommendations would help improve the security of Lawful Intercept.

4.1 Recommendations for Updating the User-based Security Model for SNMPv3

4.1.1 Make Authentication Errors Less Verbose

The User-based Security Model for SNMPv3 provides too much information in response to authentication failures, allowing attackers to differentiate between bad usernames and bad passwords. Most authentication systems have moved away from providing this level of verbosity because it is only marginally valuable to legitimate users and it is helpful to attackers.

4.1.2 Make Engine Values Harder to Guess and Share

The use of Engine values within the User-based Security Model for SNMPv3 was intended to prevent previous requests from being replayed. This feature also serves to make packet spoofing of SNMPv3 transactions more complicated by requiring an interactive transaction between client and server. However, these values do not provide complete anti-spoofing protection. The Engine values are theoretically subject to being guessed, and they can be shared between clients, a fact which enables attackers with SNMP access to subvert user group access control lists.

TCP implementations prevent spoofing by employing unpredictable sequence numbers. SYN cookies are a particularly good way to generate those sequence numbers. They combine the server’s concept of time and its own identity with the identity of the client and a secret key. They cannot be easily guessed or shared between clients, which prevents spoofing, but they also need not be cached by the server, which prevents denial of service attacks. [23] If the SNMPv3 User-based Security Model were modified to use a cookie that is similar to a SYN cookie, instead of or in addition to the Engine values, this would effectively eliminate packet spoofing as a concern for authenticated SNMPv3 transactions over UDP with very little added transaction cost.

4.1.3. Send Traps or Informs for SNMPv3 Authentication Failures

The User-based Security Model for SNMPv3 is susceptible to numerous attacks against authentication. However, all of these attacks are noisy in that they generate large numbers of authentication failures before they are successful. If these authentication failures resulted in traps or informs, they could be readily monitored by SNMP management equipment and network operators would know when they were under attack. Furthermore, sending traps or informs for SNMPv3 authentication failures would be consistent with the behavior of SNMPv2, which generates traps when invalid community strings are used.

4.2 Recommendations for Updating Cisco's Lawful Intercept Architecture

4.2.1 Use a Different Port or Protocol

Some of the vulnerability of Cisco's Lawful Intercept Architecture stems from the fact that it is an application requiring very high security that was built on top of SNMP, which has historically been used for applications that do not raise the same sort of security concerns. The most common use of SNMP is to convey information about network health, which is of little value to an attacker. It was a good idea to make use of the existing SNMP code base, which was already written and tested, rather than define a completely new protocol. But, service provider networks are full of SNMP traffic that is performing various functions, and so it is difficult to limit access to this service.

It would have been better if a unique UDP port number were assigned for this high security application. Then traditional IP and port based access control lists could effectively limit access to it, as could external packet filtering devices that are not necessarily application or credential aware. SNMP can also be run over TCP, so Cisco could require TCP connections for Lawful Intercept requests. This would help to make the interface less susceptible to packet spoofing while also making it easier to filter.

4.2.2 Allow Router Administrators to Hard Code Mediation Device Addresses

It is difficult for router administrators to control where Lawful Intercept data can be sent, because the interface allows any destination to be specified, and Cisco devices do not apply outbound access control lists to packets that originate from the router. It seems unlikely that the IP addresses, port numbers, and network protocols spoken by a particular Mediation Device are going to change every time an intercept is performed. If these settings could at least be constrained to a particular range by the router administrator this would significantly reduce the risk of unauthorized intercepts. It is unlikely that this reduced flexibility would have any negative impact on legitimate intercepts.

4.2.3 Move SNMP Notification Control Out of the MIB

Users of the TAP-MIB should not be able to disable SNMP notifications regarding intercepts they are performing. This undermines the ability of the service provider to secure a reliable audit trail of surveillance activities. In our opinion, control over intercept notifications should have been placed in the router configuration rather than in the MIB.

However, there is a tension between this desire and the need to prevent the service provider's network administrators from monitoring the activities of the Law Enforcement

Agency. It is difficult, in practice, to prevent a determined network administrator with the ability to access and reconfigure routers and switches from tracking the use of a protocol on the network, particularly if encryption is not being used. However, if a rouge administrator could configure an unauthorized destination address to receive SNMP notifications regarding the TAP-MIB, this would be a particularly easy way to monitor the use of Lawful Intercept.

One way to address this concern would be to specify destination addresses for intercept notifications in both the MIB and the device configuration. The device could enforce that at least one of the addresses it is configured to notify is included as a notification address in the Interception Request. This way the Interception Request must specify a destination for notifications that the network administrator has agreed to, but conversely the network administrator cannot force the device to send notifications to an address the Lawful Intercept Administration doesn't agree to.

4.3 Recommendations for Deployment

Service providers have many different strategies for protecting management traffic on their networks. One option is total out-of-band management. This approach essentially amounts to concluding that router and switch management interfaces cannot be protected from attack and a shell must be constructed around them. A parallel network infrastructure is created along side the provider's production network. All of the provider's routers are connected to this parallel network, and access lists are placed on the other interfaces of each router that block all inbound management traffic. Access to these out of band management networks is very tightly controlled.

While this approach is effective for some providers, it is impractical for others. While all prudent network operators have some backup mechanism for getting access to a router in the event that the primary network has failed, building a backup network that is designed to operate 24/7 and carry all remote management traffic can be very expensive. Also, as a network becomes more and more complicated with a larger management staff, the provider's ability to absolutely control everything on an out-of-band management network may come into question.

It is therefore necessary to consider how to deploy Lawful Intercept safely, no matter how tightly controlled a provider thinks their management networks are. Below is a set of basic recommendations based on the analysis in this paper.

1. Use Access Control Lists
   - Limit SNMP and other management access to routers and switches from unauthorized source IP addresses.
   - Specify an Access Control List with the IP address of the Mediation Device and apply it to the SNMPv3 User Group that has access to the TAP-MIB.

2. Lock down SNMPv3
   - Make sure your Cisco devices are not running a version of IOS that is vulnerable to the SNMP Password Vulnerability (CVE-2008-0960)
   - Manually configure a unique, random SNMPv3 EngineID.

    - Select usernames and passwords for TAP-MIB access that are difficult to guess and carefully control access to them.
    - Configure an IP address to receive SNMP notifications from Intercept devices and monitor it for suspicious activity.

3. Use encryption
    - Establish an IPsec tunnel between the Intercept Access Point and the Mediation Device and drop any unencrypted traffic between those addresses.
    - Always manage Cisco devices using SSH and other encrypted management protocols.

4. Harden your network devices against attack
    - Consult public hardening guides for Cisco routers and switches that instruct administrators on how to protect them against attack.
    - Establish processes to effectively control access to network administration access credentials and interfaces.
    - Protect the Mediation Device from unauthorized physical access as well as network intrusions.
    - Audit the security of every computer on the authorized management network and make sure they are hardened against attack.

5 Conclusions

Cisco and the IETF did the right thing when they decided to publish their architecture for lawful intercept so that we can study it and peer-review it. The Internet is a product of the academic world and it brings with it some of the values of that world, including a principle of inclusiveness in technical standards making and an architecture that is open to examination and peer review. Nearly every communications protocol that is widely used on the Internet is published in an IETF RFC that is available for anyone to read and understand. If a surveillance capability were quietly added into the core of the Internet and an attempt was made to keep it a secret, in some respects this would be antithetical to the overarching philosophy upon which the Internet was built.

As government surveillance is obviously a subject of significant public interest, the value of being able to discuss it in the open cannot be overstated. In many respects, the technical operation of these systems, in particular how they go about minimizing the information that they collect and how they prevent unauthorized access, is as much a part of the framework of checks and balances that protects the privacy rights of citizens, as the legal procedures that authorize their use. Providing open, public access to the architectural design of these systems and the procedures used to operate them enables independent experts to review their suitability to task and allows individual citizens to see and understand how surveillance is being implemented.

Although we identify some weaknesses in Cisco's Architecture for Lawful Intercept in this paper, at the end of the day, we can have greater confidence in Cisco's approach because it's a known quantity and because it has had the benefit of analysis like this. There is a lot of technology out there that the public depends on that is more difficult to peer-review, and in those shadows our greatest vulnerabilities lie.

The author would like to thank Cisco their feedback and thoughtful consideration of the issues raised by this paper. The author would also like to thank a number of individuals who provided constructive feedback that contributed to this paper: Greg Conti, David Dewey, Nick Levay, Felix Lindner, John McDonald, Jeremy Mineweaser, and Terrence Peters.

Tom Cross *is a vulnerability researcher with IBM's X-Force, but the opinions he expresses herein are his own and not those of his employer. His research interests include protocol design, application software vulnerabilities, reverse engineering, and Voice over IP. He has a BS in Computer Engineering from the Georgia Institute of Technology. Contact him at tcross@us.ibm.com.*

[1] Steven M. Bellovin, Matt Blaze, Whitfield Diffie, Susan Landau, Peter G. Neumann, and Jennifer Rexford, "Risking Communications Security: Potential Hazards of the Protect America Act," *IEEE Security & Privacy*, vol. 6, no. 1, January/February 2008, pp. 24-33.

[2] Internet Architecture Board, "IETF Policy on Wiretapping," RFC2804; http://tools.ietf.org/html/rfc2804

[3] F. Baker, B. Foster, C. Sharp, "Cisco Architecture for Lawful Intercept in IP Networks," RFC3924; http://tools.ietf.org/html/rfc3924

[4] F. Baker, "Cisco Lawful Intercept Control MIB," draft-baker-slem-mib-00; http://tools.ietf.org/html/draft-baker-slem-mib-00

[5] U. Blumenthal and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)," RFC3414; http://tools.ietf.org/html/rfc3414

[6] Common Vulnerabilities and Exposures, "CVE-2008-0960"; http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2008-0960

[7] FX, "Perception of Vulnerabilities," Aug 27, 2008; http://www.phenoelit.net/lablog/paradigms/Perception_of_Vulnerabilities.sl

[8] Hardaker, "SECURITY RELEASE: Multiple Net-SNMP Versions Released," June 9, 2008; http://sourceforge.net/forum/forum.php?forum_id=833770

[9] Cisco Systems, "How To Copy Configurations To and From Cisco Devices Using SNMP," Nov 1, 2005;

http://www.cisco.com/en/US/tech/tk648/tk362/technologies_configuration_example09186a0080094aa6.shtml

[10] Cisco Systems, “Cisco Security Advisory: SNMP Version 3 Authentication Vulnerabilities,” July 1, 2008; http://www.cisco.com/en/US/products/products_security_advisory09186a00809ac83b.shtml

[11] Cisco Systems, “C10K Lawful Intercept Configuration Guide Book”; http://www.cisco.com/en/US/docs/routers/10000/10008/feature/guides/lawful_intercept/10LIconf.html

[12] Cisco Systems, “Cisco IOS Release 12.3(7)XI2,” 2004; http://www.cisco.com/en/US/docs/routers/10000/release_notes/123xi/1237xi2.html

[13] Cisco Systems, “Cisco 10000 Series Routers,” Retrieved on November 4th, 2008; http://www.cisco.com/en/US/products/hw/routers/ps133/

[14] Cisco Systems, “KPN Delivers IP-VPN Services Over MPLS to Dutch Enterprise and SMB Customers Using Cisco Solutions,” May 21, 2002; http://newsroom.cisco.com/dlls/prod_052102.html

[15] Cisco Systems, “Individual SNMP Trap Support,” 2006; http://www.cisco.com/en/US/docs/ios/12_1t/12_1t3/feature/guide/dtitraps.html

[16] Cisco Systems, “snmp-server engineID local,” Cisco IOS Network Management Command Reference; http://www.cisco.com/en/US/docs/ios/netmgmt/command/reference/nm_14.html#wp1011412

[17] Cisco Systems, “Specifying SNMP-Server Group Names,” Configuring SNMP Support; http://www.cisco.com/en/US/docs/ios/netmgmt/configuration/guide/nm_cfg_snmp_sup_ps6441_TSD_Products_Configuration_Guide_Chapter.html#wp1026796

[18] Cisco Systems, “Cisco Applied Mitigation Bulletin: Identifying and Mitigating Exploitation of the SNMP Version 3 Authentication Vulnerabilities,” June 11, 2008; http://www.cisco.com/en/US/products/products_applied_mitigation_bulletin09186a00809adfc8.html

[19] Cisco Systems, “Cisco Service Independent Intercept Architecture Version 2.0,” 2006; http://www.cisco.com/technologies/SII/SII_ver2.pdf

[20] Cisco Systems, “Lawful Intercept on Cisco 12000 Series Router, ISE Line Cards,” 2006; http://www.cisco.com/en/US/docs/ios/12_0s/feature/guide/lawf_int.html

[21] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and Schooler, E. “Sip: Session Initiation Protocol,” RFC 3261, http://tools.ietf.org/html/rfc3261

[22] Blake, D. “Architecture & Design: Network Protocols,” *Dr. Dobb’s*, October 30, 2008; http://www.drdobbs.com/architect/211800253

[23] D.J. Bernstein. SYN cookies. 2008; http://cr.yp.to/syncookies.html