OPTIMAL AVAILABILITY & SECURITY FOR VOICE OVER CABLE NETWORKS

Chun K. Chan, Andrew R. McGee, Martin J. Glapa, and Uma Chandrashekhar
Bell Laboratories, Lucent Technologies

Abstract

The Lucent Bell Labs Security model has now become the foundation of the newly ratified ITU-T Recommendation X.805 "Security Architecture for Systems Providing End-to-End Communications." The X.805 standard was developed as the framework for the architecture and dimensions in achieving end-to-end security of distributed applications. In this paper, we introduce the X.805 standard and describe how it can be applied to the PacketCable™ Security Specification (PKT-SP-SEC-107-021127) and the DOCSIS BPI+ specification (SP-BPI+-I10-030730) for Voice over Cable (VoC) networks. We identify areas of conformance and gaps that the current PacketCable standards have with respect to the X.805 Security Model and examine the effect on end-to-end availability of VoC networks. The PacketCable™ reliability models (PKT-TR-VoIPAR-V01-001128) are generalized to include downtime due to security vulnerabilities and attacks. Our analysis shows that the traditional reliability models produce results that are optimistic if we do not consider both availability and security within a network dependability framework. The X.805 standard can be used to augment these models to provide optimal availability and security for Voice over Cable networks.

INTRODUCTION

The advent of next generation IP networks carrying converged voice and data traffic obliterates the inherent, built-in security of traditional telecommunications networks with their unintelligent end-user devices and out-of-band signaling and management networks. Now powerful, intelligent devices under end-user control are potentially able to access signaling and network management information. Hobbyists have hacked into cable modem hardware, tricking it to accept custom code. As a result, the end user has complete control of the cable modem and can surmount the bandwidth imposed by the service provider [1]. In Voice over Cable network, it is conceivable that hackers can also modify and gain control of the MTA. Rolling out VoIP services over cable automatically inherits some of the vulnerabilities associated with VoIP. A recent CERT advisory states that a number of vulnerabilities have been discovered in various implementations of the H.323 protocol [2] [3]. Even though Voice over Cable does not use H.323, we expect similar vulnerable implementations of relatively new protocols.

Given that a Voice over Cable network is vulnerable, cable operators can benefit from a comprehensive, end-to-end security framework to guide their network planning and the ongoing security assessments performed against their networks. The recently ratified ITU-T Recommendation X.805 "Security Architecture for Systems Providing End-to-End Communications" [4] was developed to provide such a framework.

VOICE OVER CABLE NETWORK

Figure 1 represents a simplified functional view of the PacketCable™ based VoIP network architecture. VoIP builds on a DOCSIS high-speed data infrastructure. The CMTS (Cable Modem Termination System) provides DOCSIS IP connectivity on the RF

based cable network with a managed data network or the Internet. For VoIP, DOCSIS is used to transport IP packets containing signaling and bearer (voice) packets between an MTA device at the subscriber location and various network elements. The MTA provides the telephony termination/origination point. The MTA can either be embedded with a cable modem (E-MTA) or a standalone device that connects to a cable modem (S-MTA). E-MTAs can take advantage of a Dynamic Quality of Service (DQoS) feature to provide priority for voice traffic, whereas an S-MTA cannot.

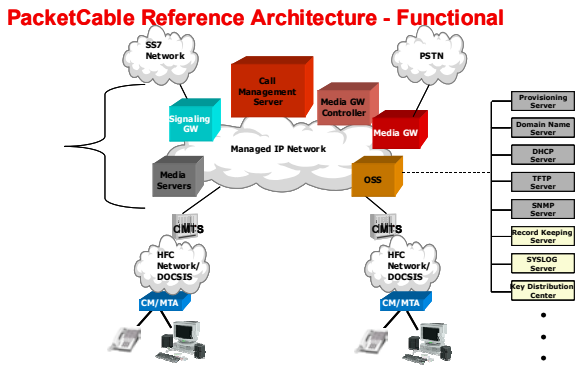

Figure 1: Functional PacketCable™ Reference Architecture

A Call Management Server (CMS) provides subscriber side call processing functions such as origination, tear-down and Class 5 switch features. A Media Gateway Controller (MGC) is a softswitch function that provides PSTN or trunk side call processing control of the Media Gateway (MG). The MG provides circuit-packet conversions for connectivity to the PSTN. Other various servers provide SS7 interfaces, announcements, voice mail and back office systems.

This type of network can be implemented in a single cable serving area or it can be implemented to span multiple cable serving areas across geographic regions as well as across multiple MSOs. The network can be distributed where softswitches and gateways and other elements are scattered across multiple locations, or centralized where the softswitch, gateways and other elements are collocated together. Vendor products can implement single discrete functions of the reference architecture or can integrate several functions into a single product.

Given all these subscriber, network and product variables, coupled with vulnerabilities mentioned earlier, security/reliability in this network presents challenges to be addressed by a comprehensive, end-to-end security framework such as ITU-T Recommendation X.805.

ITU-T RECOMMENDATION X.805

The advent of next generation cable IP networks carrying converged voice and data traffic obliterates the inherent, built-in security provided by traditional telecommunications networks. In modern IP/cable networks we now have the situation where numerous, powerful, intelligent end-user devices that can be used to launch network attacks are attached to cable networks. The signaling/control and management information is carried in-band with user information thereby making it susceptible to attack as well. Network operations or management security is often neglected when MSO network security is being considered and frequently provides a back-door entry into MSO networks. Since the insider threat represents a potential for significant financial loss, this situation is a recipe for disaster. The X.805 Security Architecture provides a structured framework that forces the consideration of all these factors to provide comprehensive, end-to-end network security.

The X.805 Security Architecture defines the framework for the architecture and dimensions in achieving end-to-end security of distributed applications. The general principles and definitions apply to all applications, even though details such as threats and vulnerabilities and the measures to counter or prevent them vary based upon the needs of the application [5]. How each standard fits together in the end-to-end security picture emanates from X.805. ITU-T Recommendation X.805 also forms the foundation for the proposed ISO/IEC 18028 standard "Information technology - Security techniques - Network Security - Part 2: Network Security Architecture," which has recently completed Committee Draft balloting in preparation to becoming an international standard.

We provide a brief description of the X.805 Security Architecture before demonstrating how it can be applied to Voice over Cable networks.

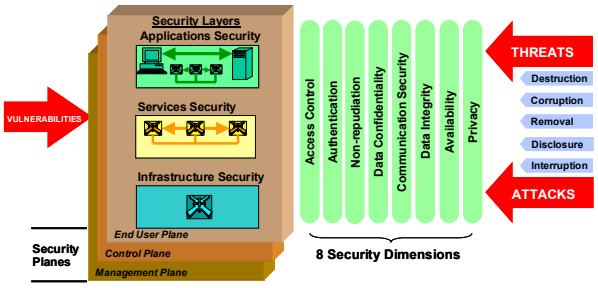

Figure 2: ITU-T Recommendation X.805 "Security Architecture for Systems Providing End-to End Communications"

The X.805 Security Architecture was developed as part of the ITU-T X.800 series of recommendations [6] to provide a methodical, organized way of addressing the five threats to telecommunications networks. The X.800 series identifies these threats as:

- *Destruction* of information and/or other resources,
- *Corruption* or modification of information,
- *Removal*, theft, or loss of information and/or other resources,
- *Disclosure* of information, and
- *Interruption* of services.

Figure 2 depicts the X.805 Security Architecture, which provides a systematic way of countering these five threats for large, complex networks such as today's MSO networks. X.805 provides a comprehensive, multi-layered, end-to-end view of network security across eight security dimensions. The X.805 standard defines a hierarchy of network equipment and facility groupings into three Security Layers: (1) the Infrastructure Security Layer, (2) the Services Security Layer, and (3) the Applications Security Layer.

- The *Infrastructure Security Layer* consists of the basic building blocks used to build telecommunications networks, services and applications, and consists of individual communication links and network elements including their underlying hardware and software platforms. Examples include the cable modem, CMTS, CMS, Signaling/Media Gateways, Media Gateway Controllers, and Media Servers depicted in Figure 1.
- The *Services Security Layer* consists of services that customers/end-users receive from networks. Example services range from basic connectivity and transport (e.g., Internet access) to service enablers (e.g., authentication, authorization, and accounting – AAA services) to value-added services such as VPN, VoIP, and Voice over Cable services.
- The *Applications Security Layer* focuses on network-based applications that are accessed by customers/end-users. These applications are enabled by network

services and are characterized by the end-user interacting with remote hardware or software in order to access information or perform a transaction. Example network-based applications include basic applications such as file transport (e.g., FTP) and web browsing, fundamental applications such as directory assistance (e.g., 411) and e-mail, as well as high-end applications such as e-commerce, network-based training, and video collaboration.

These Security Layers provide comprehensive, end-to-end security solutions and identify where security must be addressed in products and solutions because each layer may be exposed to different types of threats and attacks. For example, a Denial of Service (DoS) attack can be performed at the Infrastructure Layer by flooding a router's physical port with bogus packets, thus preventing or impeding the transmission of legitimate traffic. A DoS attack can also be performed at the Services or Applications Layer by deleting user account information, thus preventing legitimate users from accessing the service or application. One can readily see that components of Infrastructure Security, Services Security, and Applications Security must be addressed in order to provide a comprehensive, end-to-end network security solution, and that different counter-measures must be applied at each Security Layer.

Three types of activities are performed on any network, which are represented by the three Security Planes: (1) the Management Plane, (2) the Control Plane, and (3) the End-User Plane. Different security vulnerabilities may exist in each of these planes – in fact each of these planes might be implemented by separate networks for a given network or service architecture. Each Security Plane along with the three layers must be secured in order to provide an effective security posture.

The eight Security Dimensions contained in recommendation X.805 represent classes of actions that can be taken, or technologies that can be deployed, to counter the unique threats and potential attacks present at each Security Layer and Plane:

- *Access Control* is concerned with providing authorized access to network resources.
- *Authentication* is concerned with confirming the identity of communicating parties.
- *Non-repudiation* is concerned with maintaining an audit trail, so that the origin of data or the cause of an event or action cannot be denied.
- *Data Confidentiality* is concerned with protecting data from unauthorized disclosure.
- *Communication Security* is concerned with ensuring that information only flows between authorized end-points without being diverted or intercepted.
- *Data Integrity* is concerned with maintaining the correctness or accuracy of data and protecting against unauthorized modification, deletion, creation, and replication.
- *Availability* is concerned with ensuring that there is no denial of authorized access to network elements, stored information, information flows, services, and applications.
- *Privacy* is concerned with protecting information that might be derived from the observation of network activities.

Table 1 indicates how the Security Dimensions relate to the X.800 threats described previously; the cells marked with 'Y' indicate the Security Dimensions that are applicable to each of the five threats. In particular, through this mapping, we can begin to identify the right security mechanisms needed to thwart potential threats.

Security Dimension	X.800 Security Threat				
	Destruction	Corruption	Removal	Disclosure	Interruption
Access Control	Y	Y	Y	Y	
Authentication			Y	Y	
Non-repudiation	Y	Y	Y	Y	Y
Data Confidentiality			Y	Y	
Communication Security			Y	Y	
Data Integrity	Y	Y			
Availability	Y				Y
Privacy				Y	

Table 1: Applying Security Dimensions to Security Threats

The X.805 Security Architecture can also be addressed in a modular form, as illustrated in Figure 3, to provide a systematic, methodical approach to network security. Figure 3 shows the intersection of a Security Layer with a Security Plane. This represents a unique perspective for consideration of the eight Security Dimensions and can be considered a component, or module, of end-to-end network security. Each of the nine modules in Figure 3 combines the eight Security Dimensions that are applied to each security perspective. The Security Dimensions of different modules have different objectives and consequently comprise different comprehensive sets of security measures. The tabular form gives a convenient way of describing the objectives of the Security Dimensions for each module.

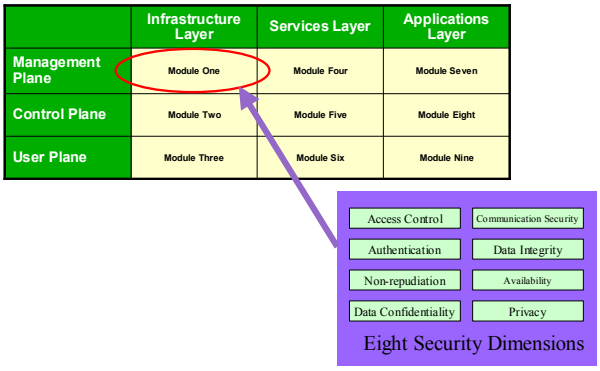

Figure 3: Modular form of X.805 Security Architecture

APPLYING X.805 TO VOICE OVER CABLE

CableLabs® has developed a security specification for providing security to VoIP communications over the PacketCable™ reference architecture described above. The PacketCable Security Specification [7] was defined to provide confidentiality to user information flows (voice and data) across the PacketCable network and to protect cable MSOs against theft of service. The PacketCable Security Specification defines the security architecture, protocols, algorithms, functional requirements and technological requirements that force any user with the intent to steal or disrupt network services to spend an unreasonable amount of money or time to do so.

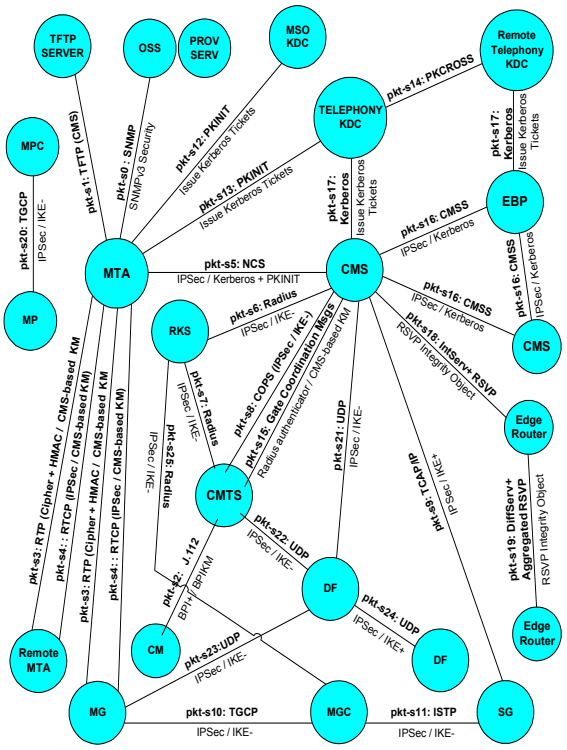

Figure 4: PacketCable™ Security Architecture

Figure 4 depicts the security architecture specified by the PacketCable Security Specification, which provides device authentication and authorization as well as encryption for the PacketCable network. The PacketCable Security Specification relies on DOCSIS 1.1 [8] and BPI+ [9] to secure the information flow across the HFC portion of the PacketCable network. The PacketCable Security Specification extends encryption to the Call Management Server (CMS), Media Gateway (MG), Signaling Gateway (SG), and remote Multimedia Terminal Adapter (MTA).

In order to provide confidentiality for user information flows and protect against theft of service, the PacketCable Security Specification contains detailed requirements for encryption algorithms as well as authentication algorithms to be used in the PacketCable network. In summary, the PacketCable Security Specification provides:

- Confidentiality of the user voice/data streams across the PacketCable network,
- Secure bearer, signaling and management channels across the PacketCable network,
- Protection against CPE cloning/tampering/uncapping,
- Protection against identity theft.

The X.805 security framework can be used to augment the Packet Cable Security Specification to provide comprehensive end-to-end security by including additional portions of the PacketCable network architecture and the X.805 Security Dimensions, Layers and Planes. When the scope of the security analysis is extended to include the end-to-end PacketCable network architecture and the entire X.805 security framework, vulnerabilities are identified in every layer, plane and dimension, even when the analysis is limited to the transport of packetized voice across the MSO network. For example, cable MSO networks must be protected against unauthorized access achieved by bridging unprotected user networks (e.g., WiFi networks) into the PacketCable environment.

Protecting the MSO network against Denial of Service attacks is another area where X.805 augments the PacketCable Security Specification. Denial of Service attacks are attacks on the Availability security dimension and are probably the most widely publicized type of security vulnerability. Denial of Service can be achieved in many different ways including: (1) an unauthorized user logging in as a system administrator causing a critical network element to crash, (2) deletion of user account information thereby preventing authorized users from accessing a service, (3) flood attacks like "smurf" that consume network resources to the point that no one can access it, (4) viruses and worms such as "Code Red" and "NIMDA" that exploit system vulnerabilities to gain access to vulnerable machines and then propagate themselves to other vulnerable hosts, which also results in the consumption of network resources.

The emphasis placed on Access Control and Authentication by the PacketCable Security Specification protects against Denial of Service attacks accomplished via unauthorized access to network elements, with the exception of End-User devices, which are considered out of scope.

Cable MSOs can use the X.805 security architecture to identify mechanisms that can be used to augment the PacketCable Security Specification to protect against the additional types of DoS attacks. As evidenced by the recent Code Red attack's ability to cripple CMTS devices throughout the world [10], every PacketCable network element (CMTS, MG, SG, CMS) as well as the Operations Support System servers, the back-office servers, etc. are potentially vulnerable to flood attacks and network worms. X.805 also

indicates that cable MSOs must also develop mechanisms to address Denial of Service attacks achieved by attacking the user information, etc., which is critical to the Voice over Cable service.

Availability Security Dimension		
X.805 Security Plane	X.805 Security Layer	
	Infrastructure	Services
End-User	Not Applicable	Missing
Control	Not Applicable	Missing
Management	Incomplete	Missing

Table 2: PacketCable Coverage for Availability Security Dimension

The Privacy Security Dimension is another example of how X.805 augments the PacketCable Security Specification to provide comprehensive end-to-end security. This dimension is concerned with protecting information about activities that take place on the network. For the Voice over Cable service, the source and destination of a communication flow would be an example of this type of information. For example, it may be important to protect the fact that two parties are communicating with each other over and above the actual contents of the communication. The PacketCable Security Specification utilizes the IPSec ESP protocol in transport mode [11], which does not encrypt the original IP packet header. Therefore, the Privacy Security Dimension is not addressed by IPSec ESP transport mode per se. The NCS messages that contain dialed numbers and other customer information are carried as IP packet payloads which are encrypted via IPSec ESP transport mode; however, once the call is established, the IP addresses of the communicating end-points are visible. The tunnel mode for the IPSec ESP protocol would provide complete coverage of the Privacy Security Dimension for these messages.

Privacy Security Dimension		
X.805 Security Plane	X.805 Security Layer	
	Infrastructure	Services
End-User	Incomplete	Incomplete
Control	Not Applicable	Missing
Management	Missing	Missing

Table 3: PacketCable Coverage for Privacy Security Dimension

This section has provided some key results of applying X.805 to a portion of the PacketCable network architecture in order to demonstrate how the PacketCable Security Specification can be augmented to achieve optimum security for the Voice over Cable service. A complete analysis of the end-to-end Voice over Cable network architecture utilizing X.805 has produced comparable results for the remainder of the VoC network architecture and remaining X.805 Security Layers, Planes and Dimensions.

IMPACT OF SECURITY VULNERABILITIES & ATTACKS ON END-TO-END AVAILABILITY

The X.805 analysis in the preceding section briefly addresses the key security challenges that need to be addressed when looking at the eight dimensions of network security. When we look at supporting key applications on a cable infrastructure such as voice we need to think about availability. In this section, we focus on the availability dimension. The standard reference for Voice over Cable availability is the Cable Labs specification, PKT-TR-VoIPAR-V01-001128 [12]. We begin with a brief review of the PacketCable™ VoIP availability allocation process.

The PacketCable™ VoIP models allocate availability budgets to network elements so that the end-to-end availability of a voice-over-packet network is the same as that of the

PSTN. For instance, the MTA is allocated an availability budget of 99.9975%, which is equivalent to an average annual downtime of 13 minutes. An MTA is expected to be built to this specification using standard reliability engineering practices such as thermal management and component derating [13]. If all the network elements in the PacketCable™ based VoIP network architecture are built to meet their respective availability budgets, then the end-to-end VoIP availability is expected to be the same as that of the PSTN. Even if all these budgets are met, we argue that it is unlikely that the end-to-end goal of PSTN availability will be met because the PacketCable™ VoIP availability/reliability models do not include downtime due to security vulnerabilities and threats. Theft of services such as MTA tampering does not have a direct impact of end-to-end availability, so such vulnerabilities and threats are not modeled here.

Denial-of-service attacks such as Code Red and NIMDA have brought down CMTS devices [10]. According to PKT-TR-VoIPAR-V01-001128, the CMTS is allocated a downtime of 10 minutes per year. To meet this downtime budget requires redundant hardware, which is represented by a generic parallel system shown in Figure 5. This redundant system is fault tolerant with respect to hardware and software faults; when the active fails, the standby takes over. However, if security vulnerability is present, it will be in both the active and standby software. A denial of service attack will bring down both the active and the standby subsystems. This common-mode failure [14] is pictorially represented by a DoS block in series with the redundant system, as shown in Figure 6.

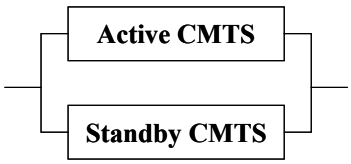

Figure 5: Active-Standby CMTS System

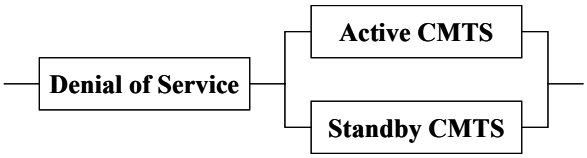

Figure 6: Denial of Service as a Common Mode Failure

To estimate the downtime due to DoS, we use the measurements from [15]. According to [15], 12,805 attacks were observed in one week on 2^{24} Internet protocol (IP) addresses. Since these 2^{24} addresses represent a theoretical maximum, we assume the number of active IP addresses to be two orders of magnitude below 2^{24}. Then, the mean attack rate per IP address is estimated to be 5×10^{-4} per hour. This implies an attack frequency of about 5 per year. Combining this with the average attack duration of 10 minutes [15], this simple model shows that DoS adds an annual downtime of 50 minutes to the CMTS system. Figure 7 shows simplified Markov models to illustrate the impact of DoS on system downtime. The state transition diagram on the left shows the scenario with no DoS. The system fails where both the active and standby units fail at the same time (duplex failure) due to hardware and/or software faults. The duplex failure rate is λ and the system restoration rate is μ. For a CMTS system that is compliant with the CableLabs specification, the system spends less than 10 minutes per year in the duplex failure state. If we include DoS, then there

will be an additional failed state labeled as the DoS state in the state transition diagram on the right in Figure 7. It is seen the system is expected to spend 50 minutes/year in the DoS state. The total system downtime is 60 minutes/year.

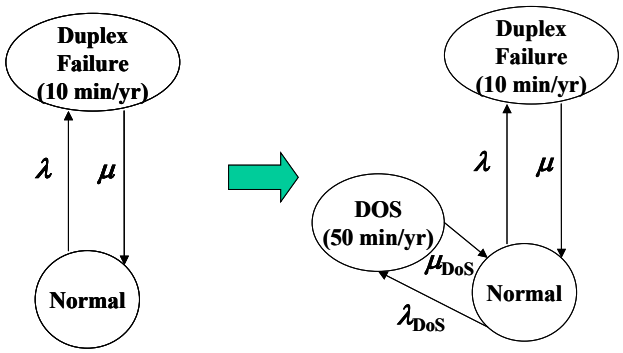

Figure 7: Simplified Markov Models Showing the Impact of DoS on System Downtime

A more detailed Markov model that takes into account of attack frequency, security vulnerability arrival rate, security vulnerability removal rate, and system restoration rate is given in [16]. The detailed model in [16] is applicable to a softswitch (such as an MGC) that uses an off-the-shelf server cluster for fault tolerance. Compared to a traditional circuit switch, which uses proprietary hardware for fault tolerance, a softswitch relies mainly on software for fault tolerance. As a result, a softswitch is more prone to security attacks [17] because it is virtually impossible to have bug-free software. An attacker could exploit well-known OS vulnerabilities to gain control of a softswitch. When the softswitch is in a compromised state, the attacker could erase critical system files so that a system re-installation is needed, then the mean time to restore service could be hours. Based on the detailed Markov model in [16], we estimate that the downtime due to this type of DoS is of the order of 100 minutes per year. This is about two orders of magnitude higher than the downtime allocated by the Cable Labs specification. As pointed out in the X.805 analysis, all PacketCable network elements (CMTS, MG, SG, CMS) as well as Operations Support System servers, back-office servers, etc. are potentially vulnerable to DoS attacks. In the remainder of this section, we show the impact on end-to-end availability if we add a DoS downtime of 100 minutes per year to the CMTS, MG, SG and CMS in the call path.

Figure 8 shows the reliability block diagram for a local on-net call between two subscribers served by the same Call Management Server (CMS). Each block contains the unavailability budget given in the Cable Labs specification. If all the network elements meet their respective unavailability or downtime budgets, then the end-to-end availability is 99.97%, which is the same as its PSTN counterpart of 150 minutes of downtime per year.

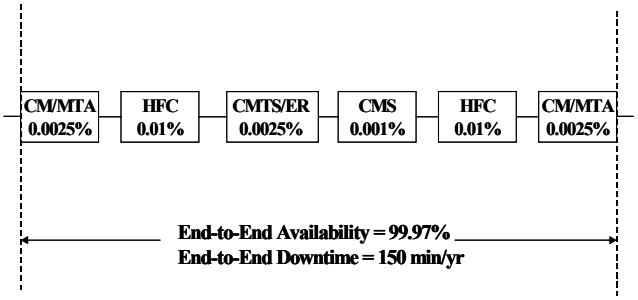

Figure 8: Local On-Net Single-Zone Availability

If we include unavailability due to DoS, then the end-to-end availability is likely to degrade from 99.97% to 99.935% (Figure 9) resulting in a downtime of 340 minutes/year. It should be noted that unavailability due to CM/MTA power outage is not included in the end-to-end calculation. If we include the downtime due to power outages [18], the degradation in end-to-end availability due to DoS is less pronounced (from 99.88% to 99.84%), because power outage alone contributes 240 minutes of downtime per year to each of the CM/MTA at both ends.

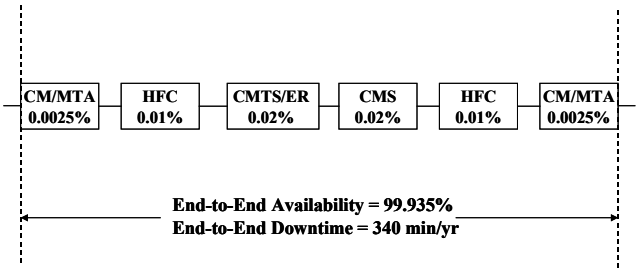

Figure 9: Local On-Net Single-Zone Availability with DoS

The impact of DoS is larger if we consider an off-net call path with more network elements that are vulnerable to DoS attacks. An off-net call is defined as a call between an endpoint on a PacketCable network and an endpoint on the PSTN. An example is given in Figure 10. This scenario shows a call where the calling and called parties are served by a CMS and a Class 5 switch with a baseline of 215 minutes of downtime per year in its normal state (without DoS attack).

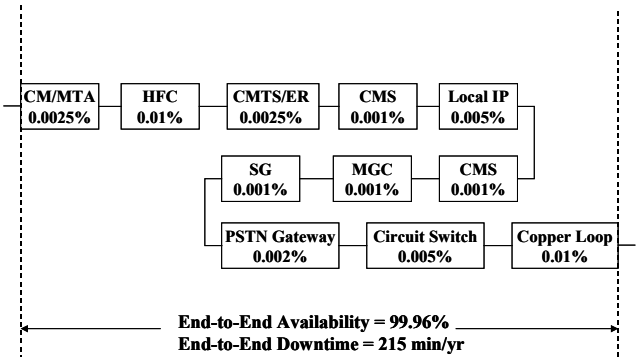

Figure 10: Local Off-Net Availability

If we include unavailability due to DoS, then the end-to-end availability is likely to degrade from 99.96% to 99.83% (Figure 11). If we also include the impact of CM/MTA power outage, the degradation in end-to-end availability is from 99.91% to 99.78%. The impact of CM/MTA power outage is less than that of the on-net scenario because we have only one CM/MTA in the call path, and the call path has more vulnerable network elements.

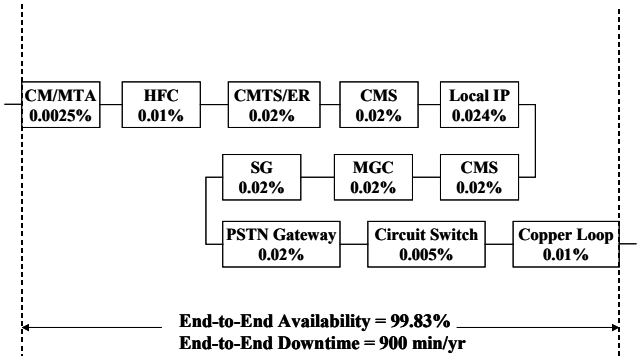

Figure 11: Local Off-Net Availability with DoS

The following table summarizes our results.

	End to End Availability (Annual Downtime)			
Scenario	Baseline	Baseline + DoS	Baseline + Power Outage	Baseline + DoS + Power Outage
Local On-Net	99.97% (150 min)	99.935% (340 min)	99.88% (630 min)	99.84% (840 min)
Local Off-Net	99.96% (215 min)	99.83% (900 min)	99.91% (470 min)	99.78% (1160 min)

Table 4: Summary of the End-to-End Availability (Downtime) Calculations

The baseline calculations in Table 4 are based on Cable Labs allocations [12]. From the end user's perspective, these numbers appear optimistic because they do not include downtimes due to DoS and power outages. For a local on-net call, the impact of power outages is larger than that of DoS, because there are two CM/MTAs in the call path and they are both affected by power outages, whereas DoS only impacts the CMTS and the CMS. For a local off-net call, the impact of DoS is larger than that of power outages, because there is only one CM/MTA in the call path and there are many network elements that are vulnerable to DoS.

Since we expect that there are more off-net calls than on-net calls, the overall (weighted) impact of DoS is larger than that of power outages. Whereas extended power outages are events that are out of the Cable MSOs' control, the impact of DoS can be mitigated by implementing software reliability engineering practices [19]. Following the downtime

allocation process in [12], the MSOs should work with their equipment vendors to allocate downtime budgets for software in addition to the existing hardware budgets. This will in turn drive the equipment vendors to improve their software development process to reduce the number of bugs and patches.

CONCLUSION

In this paper, we introduced the X.805 standard and showed how the current PacketCable™ Security Specification (PKT-SP-SEC-107-021127) and the DOCSIS BPI+ specification (SP-BPI+-I10-030730) for Voice over Cable (VoC) networks alone do not address end-to-end network security in a manner that allows cable operators to have a secure and reliable network. From the examples, we noted that in order to support VoC and other value added services, the cable operators need to have their cable network designed, maintained, and able to support an ongoing security program with controls to prevent, detect, and correct vulnerabilities resulting in maximum availability for the end-users. In particular, the controls should address the gaps noted in the security dimensions - non-repudiation, privacy, communication security, data integrity, data confidentiality, access control, availability, and authentication. By implementing these changes, and updating the models, the cable network can be designed to support VoC and the next generation of services for the end-user.

The global cost of cyber-attacks is estimated to be in the $145 billion range for 2003 alone, with 2003 also being regarded as the "worst year ever" for viruses and worms. Unfortunately, there is no end in sight to the continued onslaught of threats to network security. Clearly in today's environment, network security can no longer be treated as an afterthought and must be implemented using a continuous, systematic, methodical, end-to-end approach that has been missing until now. ITU-T Recommendation X.805 provides such an approach by providing a comprehensive, end-to-end, multi-layered view of network security across eight security dimensions.

ACRONYMS

AAA	Authentication, Authorization & Accounting
BPI+	Baseline Privacy Plus Interface
CM	Cable Modem
CMS	Call Management Server
CMTS	Cable Modem Termination System
CPE	Customer Premise Equipment
DOCSIS	Data over Cable Service Interface Specification
DQoS	Dynamic Quality of Service
DoS	Denial of Service
E-MTA	Embedded Multimedia Terminal Adapter
ER	Edge Router
ESP	Encapsulating Security Payload
FTP	File Transfer Protocol
HFC	Hybrid Fiber Coax
IP	Internet Protocol
IPSec	Internet Protocol Security
MG	Media Gateway
MGC	Media Gateway Controller
MSO	Multi-System Operator
MTA	Multimedia Terminal Adapter
NCS	Network Call Signaling
OS	Operating System
PSTN	Public Switched Telephone Network
RF	Radio Frequency
S-MTA	Standalone Multimedia Terminal Adapter
SG	Signaling Gateway
SS7	Signaling System 7
VoC	Voice over Cable
VoIP	Voice over IP
VPN	Virtual Private Network

REFERENCES

[1] K. Poulsen, “Cable Modem Hackers Conquer the Coax,” SecurityFocus, February 5, 2004.

[2] http://www.kb.cert.org/vuls/id/749342

[3] http://www.cisco.com/warp/public/707/cisco-sa-20040113-h323.shtml

[4] International Telecommunications Union, Telecommunications Standardization Sector, "Security Architecture for Systems Providing End-to-End Communications," ITU-T Rec. X.805, 2003, http://www.itu.int.

[5] International Telecommunications Union, Telecommunications Standardization Sector, "Security in Telecommunications and Information Technology," pg 1, December 2003, http://www.itu.int.

[6] International Telecommunications Union, Telecommunications Standardization Sector, "Security Architecture for Open Systems Interconnection (OSI) for CCITT Applications," ITU-T Rec. X.800, 1991, http://www.itu.int.

[7] Cable Television Laboratories, Inc., "PacketCable™ Security Specification," PKT-SP-SEC-107-021127, 2002, http://www.packetcable.com.

[8] Cable Television Laboratories, Inc., "Data-Over-Cable Service Interface Specifications DOCSIS 1.1, Radio Frequency Interface Specification," SP-RFI-v1.1-I10-030730, 2003, http://www.packetcable.com.

[9] Cable Television Laboratories, Inc., " Data-Over-Cable Service Interface Specifications DOCSIS 1.1, Baseline Privacy Plus Interface Specification," SP-BPI+-I10-030730, 2003, http://www.packetcable.com.

[10] J. T. McKelvey, “Combating Security Risks on the Cable IP Network,” IBC 2002 Conference, http://www.broadcastpapers.com/ibc2002/ibc2002.htm

[11] S. Kent and R. Atkinson, “Security Architecture for the Internet Protocol,” IETF RFC 2401, Internet Engineering Task Force, Nov. 1998, http://www.ietf.org/rfc/rfc2401.txt.

[12] Cable Television Laboratories, “VoIP Availability and Reliability Model for the PacketCable™ Architecture,” PKT-TR-VoIPAR-V01-001128, 2000.

[13] P. D. T. O’Connor, Practical Reliability Engineering, John Wiley & Sons, 1985.

[14] R. Billinton and R. N. Allan, Reliability Evaluation of Engineering Systems: Concepts and Techniques, Pitman Publishing Inc., 1983.

[15] D. Moore, G. Voelker, and S. Savage, “Inferring Internet Denial-of-Service Activities,” Proc. of the 10th USENIX Security Symposium (Washington, DC, 2001), http://www.usenix.org/events/sec01/moore/moore.pdf.

[16] C. K. Chan and H. Pant, “Reliability and Security Modeling in Upgrading Wireless Backbone Networks,” Bell Labs Technical Journal, 8(4), 39-53, 2004.

[17] Security Vulnerability in Sun Cluster 2.2, May 20, 2003, doc id 51340.

[18] Allen L, Black, James L. Spencer and Douglas S. Dorr, The Impact of Commercial Power Quality on Fiber-in-the-Loop Service Availability, Ninth Annual NFOEC Conference Proceedings, June 1993.

[19] John Viega and Gary McGraw, Building Secure Software: How to Avoid Security Problems the Right Way, Addison-Wesley, 2001.