The Internet Protocol Journal

December 1998 Volume 1, Number 3

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

You can download previous issues of IPJ in PDF format from:
www.cisco.com/ipj

From The Editor

The *Simple Network Management Protocol* (SNMP) was first standardized in 1988. It quickly became a de facto management standard, not only for Internet technologies, but for a wide range of applications. Like many early Internet protocols, the first two versions of SNMP did not include provisions for security. In 1996, two different proposals for security enhancements to SNMPv2 were put forward, with strong proponents behind each. Everyone agreed that the industry needed just *one* solution, and therefore work proceeded to incorporate the best features of the two security proposals for SNMPv2. The result is SNMPv3, and it is described in this issue by William Stallings.

As the Internet continues to grow, demand for high-speed access for residential users is increasing. Alternatives to traditional dialup service include *Digital Subscriber Line* (DSL) services, wireless solutions, and various television technologies. In this issue, we examine two aspects of Internet access using TV technologies. First, Mark Laubach gives an overview of cable modem technologies and standards, and discusses some deployment issues. In the second article, George Abe looks at the emerging digital television standards and how they could be used to provide Internet access.

The Internet lost one of its most respected pioneers when Jon Postel passed away on October 16, 1998. Jon was well-known as the Director of the *Internet Assigned Numbers Authority* (IANA) and as the editor of the *Request for Comments* (RFC) document series. Included in this issue is "I Remember IANA," a tribute to Jon Postel written by his longtime friend Vint Cerf. The remembrance has also been published as RFC 2468.

With that we have come to the end of 1998 and the end of Volume 1 of *The Internet Protocol Journal*. We wish you a pleasant holiday season and will be back with Volume 2, Number 1 in March 1999. In the meantime, please visit our Web site at **www.cisco.com/ipj**. There you will find back issues in PDF format, our Call for Papers and guidelines for authors of IPJ articles.

—Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

Security Comes to SNMP: The New SNMPv3 Proposed Internet Standards

by William Stallings

Data networks typically include bridges, routers, links into WANs, and end-user equipment from multiple vendors. Users need automated tools to help manage such configurations that are easy to install, easy to use, and don't place a great burden on the network.

This accounts for the popularity of the *Simple Network Management Protocol* (SNMP). Introduced in 1988 to provide management capability for TCP/IP-based networks, SNMP rapidly became the most widely used standardized network management tool. Virtually all vendors of network-based equipment provide SNMP.

The appeal of SNMP has indeed been its simplicity because SNMP provides a bare-bones set of functions, and it is indeed easy to implement, install, and use. And, used sensibly, it will not place undue burden on the network. Moreover, because of its simplicity, achievement of interoperability is a relatively straightforward task: SNMP modules from different vendors can be made to work together with minimal effort.

SNMP—Strengths and Weaknesses

SNMP is based on three concepts: *managers, agents,* and the *Management Information Base* (MIB). In any configuration, at least one manager node runs SNMP management software. Network devices to be managed, such as bridges, routers, servers, and workstations, are equipped with an agent software module. The agent is responsible for providing access to a local MIB of objects that reflects the resources and activity at its node. The agent also responds to manager commands to retrieve values from the MIB and to set values in the MIB. An example of an object that can be retrieved is a counter that keeps track of the number of packets sent and received over a link into the node; the manager can track this value to monitor the load at that point in the network. An example of an object that can be set is one that represents the state of a link; the manager could disable the link by setting the value of the corresponding object to the disabled state.

Such capabilities are fine for implementing a basic network-management system. To enhance this basic functionality, a new version of SNMP was introduced in 1993 and revised in 1996. SNMPv2 added bulk transfer capability and other functional extensions. However, neither SNMPv1 nor SNMPv2 offers security features. Specifically, SNMPv1/v2 can neither authenticate the source of a management message nor provide encryption. Without authentication, it is possible for nonauthorized users to exercise SNMP network management functions. It is also possible for nonauthorized users to eavesdrop on

management information as it passes from managed systems to the management system. Because of these deficiencies, many SNMPv1/v2 implementations are limited to simply a read-only capability, reducing their utility to that of a network monitor; no network control applications can be supported.

Enter SNMPv3

To correct the security deficiencies of SNMPv1/v2, SNMPv3 was issued as a set of Proposed Standards in January 1998 (Table 1). This set of documents does not provide a complete SNMP capability but rather defines an overall SNMP architecture and a set of security capabilities. These are intended to be used with the existing SNMPv2. As one of the SNMPv3 working documents puts it, "SNMPv3 is SNMPv2 plus administration and security."

Table 1: SNMPv3 RFCs

RFC Number	Title
2271	An Architecture for Describing SNMP Management Frameworks
2272	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
2273	SNMPv3 Applications
2274	User-Based Security Model for SNMPv3
2275	View-Based Access Control Model (VACM) for SNMP

SNMPv3 includes three important services: *authentication, privacy,* and *access control* (Figure 1). To deliver these services in a flexible and efficient manner, SNMPv3 introduces the concept of a *principal*, which is the entity on whose behalf services are provided or processing takes place. A principal can be an individual acting in a particular role; a set of individuals, with each acting in a particular role; an application or set of applications; or combinations thereof. In essence, a principal operates from a management station and issues SNMP commands to agent systems. The identity of the principal and the target agent together determine the security features that will be invoked, including authentication, privacy, and access control. The use of principals allows security policies to be tailored to the specific principal, agent, and information exchange, and gives human security managers considerable flexibility in assigning network authorization to users.

Figure 1: SNMPv3 Security Features

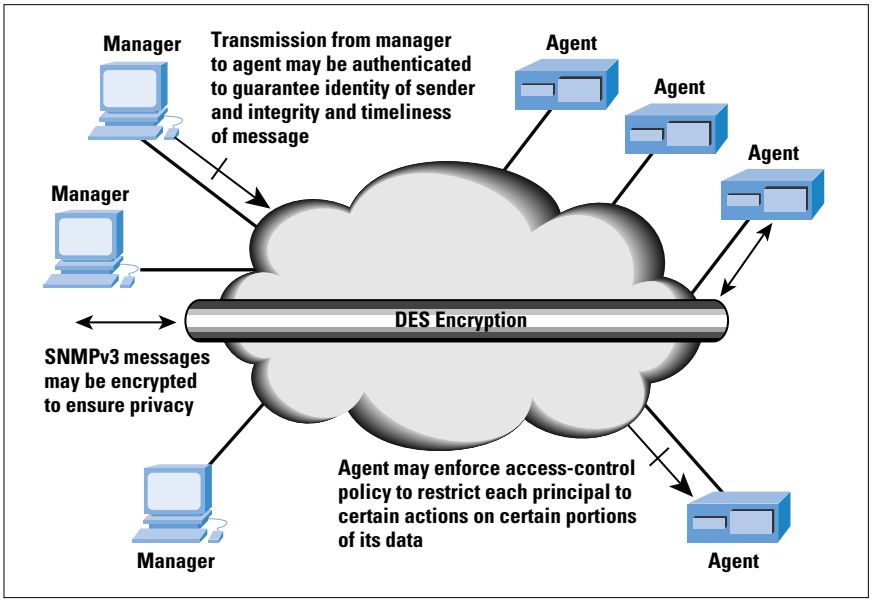

SNMPv3 is defined in a modular fashion, as shown in Figure 2. Each SNMP entity includes a single SNMP *engine*. An SNMP engine implements functions for sending and receiving messages, authenticating and encrypting/decrypting messages, and controlling access to managed objects. These functions are provided as services to one or more applications that are configured with the SNMP engine to form an SNMP *entity*. This modular architecture provides several advantages. First, the role of an SNMP entity is determined by the modules that are implemented in that entity. For example, a certain set of modules is required for an SNMP agent, whereas a different (though overlapping) set of modules is required for an SNMP manager. Second, the modular structure of the specification lends itself to defining different versions of each module. This, in turn, makes it possible to (1) define alternative or enhanced capabilities for certain aspects of SNMP without needing to go to a new version of the entire standard (for example, SNMPv4), and (2) clearly specify coexistence and transition strategies.

Figure 2: SNMP Entity (RFC 2271)

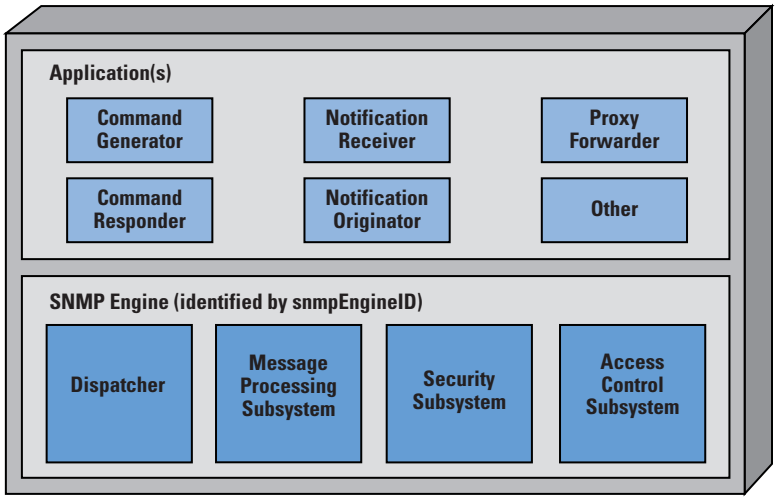

Table 2 provides a brief definition of each module.

Table 2: Components of an SNMP Entity (RFC 2271 and 2273)

Dispatcher	Allows for concurrent support of multiple versions of SNMP messages in the SNMP engine. It is responsible for (1) accepting protocol data units (PDUs) from applications for transmission over the network and delivering incoming PDUs to applications; (2) passing outgoing PDUs to the Message Processing Subsystem to prepare as messages, and passing incoming messages to the Message Processing Subsystem to extract the incoming PDUs; and (3) sending and receiving SNMP messages over the network.
Message Processing Subsystem	Responsible for preparing messages for sending and for extracting data from received messages.
Security Subsystem	Provides security services such as the authentication and privacy of messages. This subsystem potentially contains multiple Security Models.
Access Control Subsystem	Provides a set of authorization services that an application can use for checking access rights. Access control can be invoked for retrieval or modification request operations and for notification generation operations.
Command Generator	Initiates SNMP Get, GetNext, GetBulk, or Set request PDUs and processes the response to a request that it has generated.
Command Responder	Receives SNMP Get, GetNext, GetBulk, or Set request PDUs destined for the local system as indicated by the fact that the contextEngineID in the received request is equal to that of the local engine through which the request was received. The command responder application performs the appropriate protocol operation, using access control, and generates a response message to be sent to the originator of the request.
Notification Originator	Monitors a system for particular events or conditions, and generates Trap or Inform messages based on these events or conditions. A notification originator must have a mechanism for determining where to send messages, and which SNMP version and security parameters to use when sending messages.
Notification Receiver	Listens for notification messages, and generates response messages when a message containing an Inform PDU is received.
Proxy Forwarder	Forwards SNMP messages. Implementation of a proxy forwarder application is optional.

SNMPv3 Message Processing

SNMPv3 relies on the *User Datagram Protocol* (UDP) or some other transport-layer protocol to convey SNMP information. Above the UDP layer, SNMP functionality is organized into two application-level layers: a PDU processing layer and a message processing layer.

The topmost layer is the PDU processing layer. At this layer, management commands (such as Get, Set, Trap, Inform) are realized in a PDU that includes an indication of the command type and a list of variables (management objects) to which the command refers. This PDU is then passed down to the message processing layer, which adds a message header. The message header contains security-related information that may be used for authentication and privacy operations.

Figure 3 illustrates the message structure. The first five fields are generated by the message processing model on outgoing messages and processed by the message processing model on incoming messages. The next six fields show security parameters used by the security model, which is invoked by the message processing model to provide security services. Finally, the PDU, together with the contextEngineID and contextName, constitute a scoped PDU, used for PDU processing.

Figure 3: SNMPv3 Message Format with User-Based Security Model

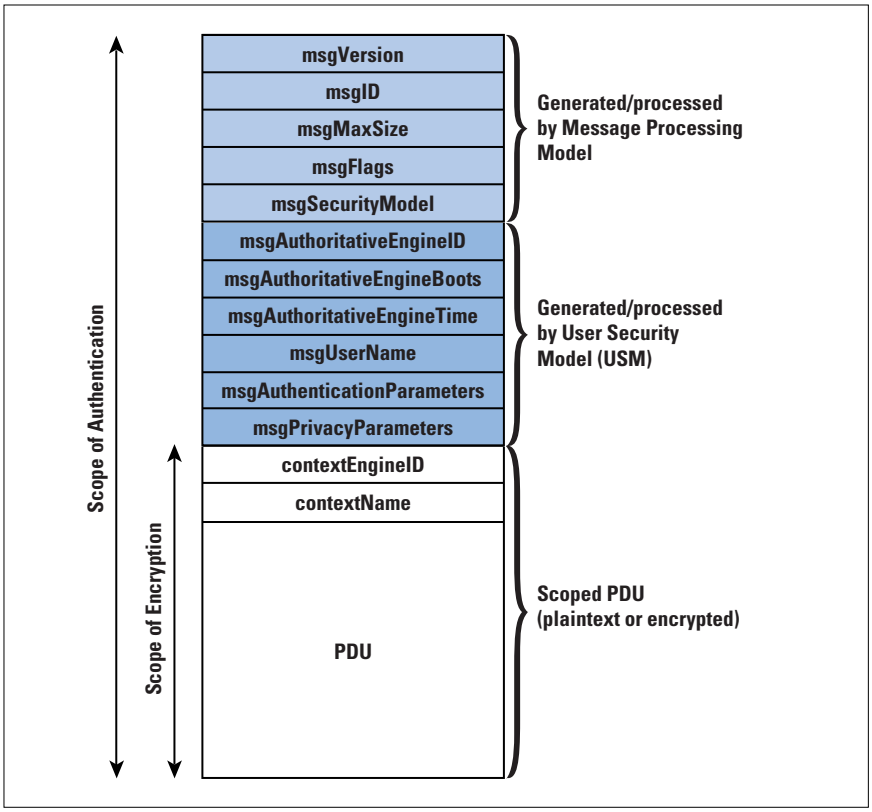

The first five fields follow:

- *msgVersion:* Set to snmpv3(3).
- *msgID:* A unique identifier used between two SNMP entities to coordinate request and response messages, and by the message processor to coordinate the processing of the message by different subsystem models within the architecture. The range of this ID is 0 through 2^{31}–1.

- *msgMaxSize:* Conveys the maximum size of a message in octets supported by the sender of the message, with a range of 484 through 2^{31}–1. This is the maximum segment size that the sender can accept from another SNMP engine (whether a response or some other message type).

- *msgFlags:* An octet string containing three flags in the least significant three bits: reportableFlag, privFlag, authFlag. If reportableFlag = 1, then a Report PDU must be returned to the sender under those conditions that can cause the generation of a Report PDU; when the flag is zero, a Report PDU may not be sent. The reportableFlag is set to 1 by the sender in all messages containing a request (Get, Set) or an Inform, and set to 0 for messages containing a Response, a Trap, or a Report PDU. The reportableFlag is a secondary aid in determining when to send a Report. It is used only in cases in which the PDU portion of the message cannot be decoded (for example, when decryption fails because of incorrect key). The privFlag and authFlag are set by the sender to indicate the security level that was applied to the message. For privFlag = 1, encryption was applied and for privFlag = 0, authentication was applied. All combinations are allowed except (privFlag = 1 AND authFlag = 0); that is, encryption without authentication is not allowed.

- *msgSecurityModel:* An identifier in the range of 0 through 2^{31}–1 that indicates which security model was used by the sender to prepare this message and, therefore, which security model must be used by the receiver to process this message. Reserved values include 1 for SNMPv1, 2 for SNMPv2c, and 3 for SNMPv3.

User-Based Security Model

The *User-Based Security Model* (USM) uses the concept of an authoritative engine. In any message transmission, one of the two entities, transmitter or receiver, is designated as the authoritative SNMP engine, according to the following rules:

- When an SNMP message contains a payload that expects a response (for example, a Get, GetNext, GetBulk, Set, or Inform PDU), then the receiver of such messages is authoritative.

- When an SNMP message contains a payload that does not expect a response (for example, an SNMPv2-Trap, Response, or Report PDU), then the sender of such a message is authoritative.

Thus, for messages sent on behalf of a Command Generator and for Inform messages from a Notification Originator, the receiver is authoritative. For messages sent on behalf of a Command Responder or for Trap messages from a Notification Originator, the sender is authoritative. This designation serves two purposes:

- The timeliness of a message is determined with respect to a clock maintained by the authoritative engine. When an authoritative engine sends a message (Trap, Response, Report), it contains the current value of its clock, so that the nonauthoritative recipient can synchronize on that clock. When a nonauthoritative engine sends a message (Get, GetNext, GetBulk, Set, Inform), it includes its current estimate of the time value at the destination, allowing the destination to assess the timeliness of the message.

- A key localization process, described later, enables a single principal to own keys stored in multiple engines; these keys are localized to the authoritative engine in such a way that the principal is responsible for a single key but avoids the security risk of storing multiple copies of the same key in a distributed network.

When an outgoing message is passed to the USM by the Message Processor, the USM fills in the security-related parameters in the message header. When an incoming message is passed to the USM by the Message Processor, the USM processes the values contained in those fields. The security-related parameters include the following:

- *msgAuthoritativeEngineID:* The snmpEngineID of the authoritative SNMP engine involved in the exchange of this message. Thus, this value refers to the source for a Trap, Response, or Report, and to the destination for a Get, GetNext, GetBulk, Set, or Inform.

- *msgAuthoritativeEngineBoots:* The snmpEngineBoots value of the authoritative SNMP engine involved in the exchange of this message. The object snmpEngineBoots is an integer in the range 0 through $2^{31}-1$ that represents the number of times that this SNMP engine has initialized or reinitialized itself since its initial configuration.

- *msgAuthoritativeEngineTime:* The snmpEngineTime value of the authoritative SNMP engine involved in the exchange of this message. The object snmpEngineTime is an integer in the 0 through $2^{31}-1$ range that represents the number of seconds since this authoritative SNMP engine last incremented the snmpEngineBoots object. Each authoritative SNMP engine is responsible for incrementing its own snmpEngineTime value once per second. A non-authoritative engine is responsible for incrementing its notion of snmpEngineTime for each remote authoritative engine with which it communicates.

- *msgUserName:* The user (principal) on whose behalf the message is being exchanged.

- *msgAuthenticationParameters:* Null if authentication is not being used for this exchange; otherwise, this is an authentication parameter. For the current definition of USM, the authentication parameter is a message authentication code generated using an algorithm referred to as HMAC.

- *msgPrivacyParameters:* Null if privacy is not being used for this exchange; otherwise, this is a privacy parameter. For the current definition of USM, the privacy parameter is a parameter used in the encryption algorithm DES.

Secret-Key Authentication

The authentication mechanism in SNMPv3 assures that a received message was, in fact, transmitted by the principal whose identifier appears as the source in the message header. In addition, this mechanism assures that the message was not altered in transit and that it was not artificially delayed or replayed.

To achieve authentication, each pair of principal and remote SNMP engines that wishes to communicate must share a secret authentication key. The sending entity provides authentication by including a message authentication code with the SNMPv3 message it is sending. This code is a function of the contents of the message, the identity of the principal and engine, the time of transmission, and a secret key that should be known only to the sender and the receiver. The secret key must initially be set up outside of SNMPv3 as a configuration function. That is, the configuration manager or network manager is responsible for distributing initial secret keys to be loaded into the databases of the various SNMP managers and agents. This can be done manually or by using some form of secure data transfer outside of SNMPv3. When the receiving entity gets the message, it uses the same secret key to calculate the message authentication code again. If the receiver's version of the code matches the value appended to the incoming message, then the receiver knows that the message can only have originated from the authorized manager, and that the message was not altered in transit. The shared secret key between sending and receiving parties must be preconfigured.

Another aspect of USM authentication is timeliness verification. USM is responsible for assuring that messages arrive within a reasonable time window to protect against message delay and replay attacks. Two functions support this service: synchronization and time-window checking.

Each authoritative engine maintains two values, snmpEngineBoots and snmpEngineTime, that keep track of the number of boots since initialization and the number of seconds since the last boot. These values are placed in outgoing messages in the fields msgAuthoritativeEngineBoots and msgAuthoritativeEngineTime. A nonauthoritative engine maintains synchronization with an authoritative engine by maintaining local copies of snmpEngineBoots and snmpEngineTime for each remote authoritative engine with which it communicates. These values are updated on receipt of an authentic message from the remote authoritative engine. Between these message updates, the nonauthoritative

engine increments the value of snmpEngineTime for the remote authoritative engine to maintain loose synchronization. These values are inserted in outgoing messages intended for that authoritative engine.

When an authoritative engine receives a message, it compares the incoming boot and time values with its own boot and time values. If the boot values match and if the incoming time value is within 150 seconds of the actual time value, then the message is declared to be within the time window and, therefore, to be a timely message.

Privacy Using Conventional Encryption

The SNMPv3 USM privacy facility enables managers and agents to encrypt messages to prevent eavesdropping by third parties. Again, manager entity and agent entity must share a secret key. When privacy is invoked between a principal and a remote engine, all traffic between them is encrypted using the *Data Encryption Standard* (DES). The sending entity encrypts the entire message using the DES algorithm and its secret key, and sends the message to the receiving entity, which decrypts it using the DES algorithm and the same secret key. Again, the two parties must be configured with the shared key.

The *cipher-block-chaining* (CBC) mode of DES is used by USM. This mode requires that an initial value (IV) be used to start the encryption process. The msgPrivacyParameters field in the message header contains a value from which the IV can be derived by both sender and receiver.

View-Based Access Control Model (VACM)

The access control facility makes it possible to configure agents to provide different levels of access to the agent's MIB to different managers. An agent entity can restrict access to its MIB for a particular manager entity in two ways. First, it can restrict access to a certain portion of its MIB. For example, an agent may restrict most manager principals to viewing performance-related statistics and allow only a single designated manager principal to view and update configuration parameters. Second, the agent can limit the operations that a principal can use on that portion of the MIB. For example, a particular manager principal could be limited to read-only access to a portion of an agent's MIB. The access control policy to be used by an agent for each manager must be preconfigured; it essentially consists of a table that details the access privileges of the various authorized managers. Unlike authentication, which is done by user, access control is done by group, where a group may be a set of multiple users.

Figure 4: VACM Flowchart

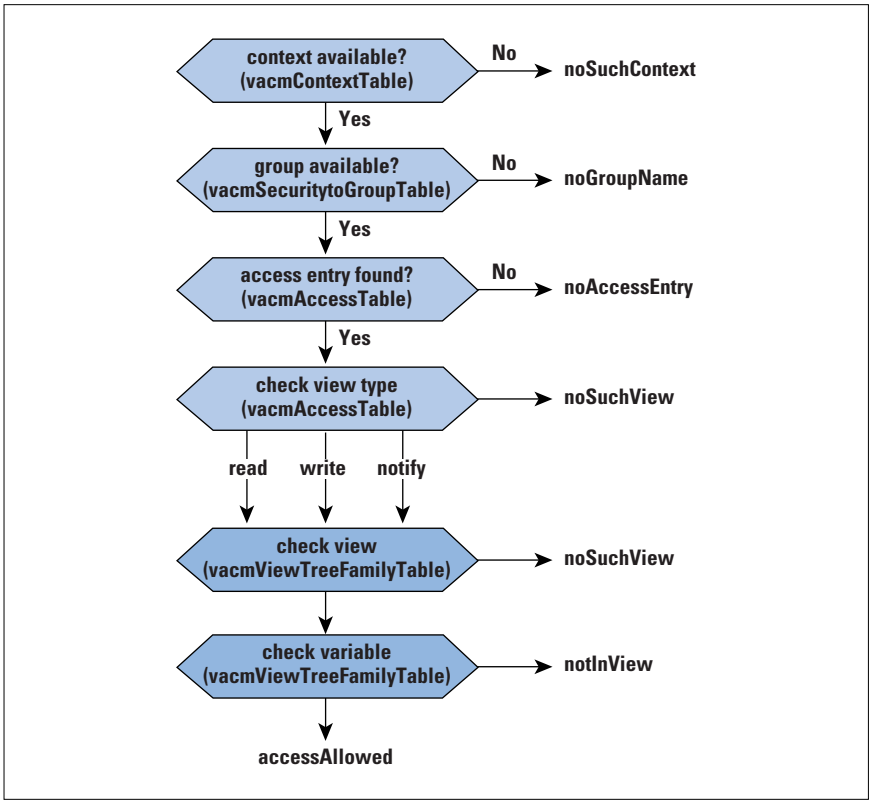

Figure 4 illustrates the overall VACM logic, which proceeds in the following steps:

1. The context name refers to a named subset of the MIB objects at an agent. VACM checks to see if there is an entry in vacmContextTable for the requested contextName. If so, then this context is known to this SNMP engine. If not, then an errorIndication of noSuchContext is returned.

2. Each principal operating under a given security model is assigned to at most one group, and access privileges are configured on a group basis. VACM checks vacmSecurityToGroupTable to determine if there is a group assigned to the requested <securityModel, securityName> pair. If so, then this principal, operating under this securityModel, is a member of a group configured at this SNMP engine. If not, then an errorIndication of noGroupName is returned.

3. VACM next consults the vacmAccessTable with groupName, contextName, securityModel, and securityLevel (indicates authentication, authentication plus privacy, or neither) as indices. If an entry is found, then an access control policy has been defined for this groupName, operating under this securityModel, at this securityLevel, for access to this contextName. If not, then an errorIndication of noAccessEntry is returned.

4. A MIB view is a structure subset of a context; it is essentially a set of managed object instances viewed as a set for access control purposes. VACM determines whether the selected vacmAccessTable entry includes reference to a MIB view of viewType (read, write, notify). If so, then this entry contains a viewName for this combination of groupName, contextName, securityModel, securityLevel, and viewType. If not, then an errorIndication of noSuchView is returned.

5. The viewName from Step 4 is used as an index into vacmViewTreeFamilyTable. If a MIB view is found, then a MIB view has been configured for this viewName. If not, then an errorIndication of noSuchView is returned.

6. VACM checks the variableName against the selected MIB view. If this variable is included in the view, then a statusInformation of accessAllowed is returned. If not, then an errorIndication of notInView is returned.

References

[0] The SNMPv3 RFCs, see Table 1 above.

[1] J. D. Case, M. Fedor, M. L. Schoffstall, and C. Davin, "Simple Network Management Protocol," RFC 1157, May 1990.

[2] Rose, M. T., *The Simple Book: An Introduction to Networking Management,* Revised Second Edition, Prentice-Hall, ISBN 0-13-451659-1, 1996.

[3] Waters, G., Editor, "User-based Security Model for SNMPv2," RFC 1910, February 1996.

[4] *ConneXions—The Interoperability Report,* Volume 10, No. 5, May 1996—Special Issue: "Network Management Today."

WILLIAM STALLINGS is a consultant, lecturer, and author of over a dozen books on data communications and computer networking. He has a PhD in computer science from M.I.T. This article is based on material in the author's latest book: *SNMP, SNMPv2, SNMPv3, and RMON 1 and 2,* Second Edition (Addison Wesley, 1998). His home in cyberspace is **`http://www.shore.net/~ws`** and he can be reached at **`ws@shore.net`**

Residential Area CATV Broadband Internet Technology: Current Status

by Mark Laubach, Com21, Inc.

Cable modem technology has entered commonplace discussion and is in the early stages of widespread deployment throughout the world. The capabilities provided by cable modems promise data bandwidth speeds far in excess of those provided by traditional telephone modem services. In North America the race is on between cable operators deploying services based on standardized cable modems and telephone companies deploying *Digital Subscriber Line* (DSL) services. Internet Service Providers (ISPs) are taking position to promote any method of delivering Internet services to and from the home and are helping to fuel the race. Initially these services will only provide higher-speed Internet access and improved access to major information services (for example, AOL). Cable modem service offerings promote higher than DSL speed to the subscriber and a promise that packet voice services will be available in 1999.

As an introduction to some of the issues surrounding cable modem technology, this article summarizes two of the standardization efforts: the IEEE 802.14 Cable TV Media Access Control and Physical Protocol working group and the North American Data Over Cable Service Interface Specification. Delivering a viable Internet service to a cable TV reached subscriber community has its own set of deployment issues that are briefly reviewed and summarized.

Background

Networks based on packet technology were first presented in 1964[1]. Since then, and through numerous evolutionary steps, the Internet as we know it today was brought into existence. Today, packets are transmitted over most any media. The next economic and technical frontier is the mass deployment of moving packets over cable television (CATV) networks for serving the Internet to every home. There are several link layer approaches for delivering IP datagrams via cable modems. The always present debate of whether to use fixed or variable length packets continues in the cable modem world. This article presents overviews of two variations of cable modem protocols: first, the concept of sending small, fixed-sized packets over the CATV plant using 53-octet *Asynchronous Transfer Mode* (ATM) cells[2], as is being defined in the public standards process of the IEEE 802.14 working group; and secondly, by sending variable-length packets (IP over Ethernet) as defined by the *Multimedia Cable Network System* (MCNS) *Data Over Cable Service Interface Specification* (DOCSIS) for the North American cable industry[3]. As widely accepted standards normally motivate industrial focus and subsequent cost reduction due to vendor competitive pressures, there is an additional drive provided by North American cable operators to get the cost of the cable modem off their books and into retail channels.

The IEEE 802.14 Cable TV MAC and PHY Protocol working group is chartered with providing a single *Media Access Control* (MAC) and multiple physical sublayer (PHY) standard for cable TV networks. The efforts of 802.14 must support IEEE 802 layer services (including Ethernet) and must also be ATM compatible.

The DOCSIS specifications are managed by CableLabs on behalf of its cable television system operator members. The project was initiated by an organization called *Multimedia Cable Network System* (MCNS) Partners, L. P., which consists of Comcast Cable Communications, Cox Communications, Tele-Communications, Inc., and Time Warner Cable. In addition to MCNS, Rogers Cablesystems Limited, MediaOne, and CableLabs have all contributed to the DOCSIS documents, as have several networking and telecommunications vendors. DOCSIS documents describe the internal and external network interfaces for a system that allows bidirectional transfer of IP traffic, between the cable system head-end and customer premises, over a cable television system[4].

The customer network interface in common use today is Ethernet 10BaseT. There is a mandate for a 10 Mbps Ethernet interface in the home. Subscriber access equipment can be a personal computer, X-Terminal, or any such device that supports the TCP/IP protocol suite. Future home interfaces from the cable modem will include the *Universal Serial Bus* (USB) and IEEE 1394 (*FireWire*).

IP Over CATV System Challenges

From an IP perspective, a CATV system almost appears to be another data link layer. However, experience gained thus far has demonstrated that the marriage of IP over CATV radio frequency (RF) channels is not as straightforward as IP over any other high-speed serial point-to-point link.

In the CATV space, the downstream channels in a cable plant (cable head-end to subscribers) is a point-to-multipoint channel. This does have very similar characteristics to transmitting over an Ethernet segment where one transmitter is being listened to by many receivers. The major difference is that baseband modulation has been replaced by a more densely modulated RF carrier with very sophisticated adaptive signal processing and *forward error correction* (FEC).

In the upstream direction (subscriber cable modems transmitting towards the head-end) the environment is many transmitters and one receiver. This introduces the need for precise scheduling of packet transmissions to achieve high utilization and precise power control so as to not overdrive the receiver or other amplifier electronics in the cable system. Since the upstream direction is like a single receiver with many antennas, the channels are much much more susceptible to interfering noise products[5, 6]. In the cable industry, we generally

call this *ingress noise*. As ingress noise is an inherent part of CATV plants, the observable impact is an unfortunate rise in the average noise floor in the upstream channel. To overcome this noise jungle, upstream modulation is not as dense as in the downstream and we have to use more effective FEC as used in the downstream. There is a further complication that there are many upstream "ports" on a fully deployed *Hybrid Fiber-Coaxial* (HFC) plant that requires matching head-end equipment ports for high-speed data[7].

To further the rub on the upstream channel use, the arcane regulations of the FCC from back in the mid 1980s mandated that upstream frequency spectrum be reserved on all cable plants, regardless of whether it was actually used. This was typically the 5–42 MHz region, leaving above 50 MHz for downstream transmissions. (Note that there are other regions available for upstream, but the overwhelming majority of cable plants only use 5–42 MHz.) This leaves precious little spectral bandwidth for upstream communications.

The existing environment for high-speed data protocols therefore provides for relatively clean bandwidth in the downstream direction, allowing for higher-speed data rate channels, while in the upstream, individual channels are of lesser data rate. However, multiple upstream channels can be used per downstream channel to get effective symmetric aggregate bandwidth. Typically, we speak of cable modem systems as providing asymmetric services (higher downstream data rate than upstream). Note though that this asymmetry closely matches what we expect initially for residential high-speed data services. That is, many more subscribers at home pulling things off the Internet via web services, than pushing data back in.

Modern modulation techniques provide for a range of data carrying capability ("baud rate"). A low order modulation rate called *Quadrature Phase Shift Keying* (QPSK) provides for two data bits per symbol encoding. *Quadrature Amplitude Modulation* (QAM) provides a lower order modulation of 16 QAM (four bits per symbol) through higher order rates of 64 QAM (six bits per symbol) and 254 QAM (eight bits per symbol). Low order modulations are more robust in higher average noise environments. Higher order modulations are least robust. Therefore, high order modulations are suitable for downstream channels due to the low noise performance, while the order of upstream channel modulation is heavily effected by noise. Typically, cable modem systems will see QPSK used for upstream channels. When the plant is very clean, noise-wise, 16 QAM may be used.

One additional challenge is that the speed of RF signals in fiber and coaxial cable is much less than the speed of light. For system deployments to be effective, the cable modem protocols must support cable modems out to a wire distance of 50 miles (80 km).

At these distances the round trip propagation delay will be on the order of 800 microseconds; which is several times the length of time it takes to transmit a 64-byte packet on the upstream channel. The IEEE and DOCSIS cable modem protocols have been engineered to overcome these propagation delays in order to increase channel utilization; that is demand-based scheduling of a slotted upstream channel coupled with precise station ranging and timing.

Another challenge is in using an IP-over Ethernet approach to providing a reliable public switched packet service to an abundance of subscribers. Traditional Ethernet networking has always relied on all the Ethernet stations being within the same administrative walls with all users sharing the same common interests. Not so with metropolitan area public access networks. Data communications must now be encrypted such that the privacy of user communications is not invaded by promiscuous neighbors. In addition, users are paying for access in this cable modem world, and any abusive behavior of users must be contained so as to not affect other users. This calls for sophisticated fairness scheduling in the head-end systems and the use of comprehensive cryptological and packet filtering techniques. It is all very complicated both to create, and to manage. Each standard has its own approach for dealing with these issues.

Where IP over CATV appeared to be fundamentally similar to Ethernet when the industry first started out, in reality it is not. High-speed cable data networking, as demonstrated by the work output from various standards activities, is fundamentally a new approach to what at first appeared to be similar old problems. It's not ALOHA anymore[8], nor is it your grandfather's Ethernet[9, 10].

IEEE 802.14 Cable TV MAC and PHY Protocol Working Group

Let's briefly examine the first comprehensive standard activity created to address the current emerging world of high-speed cable data systems. In November 1994, the IEEE 802.14 CATV MAC and PHY Protocol working group met for the first time as an approved project within the 802 standards committee. Previous work had been done in 1993 through 1994 in the 802.catv study group in preparation for formal IEEE 802 project approval. The *Project Authorization Request* (PAR) charter of the group specifies that it will standardize a single MAC layer protocol and multiple PHY layer protocols for two-way HFC networks. Consistent with the IEEE LAN/MAN 802 Reference Model[11], 802.14 is producing a solution that supports the 802 protocol stack while at the same time supporting ATM in an ATM-compatible manner.

The general 802.14 requirements include:

- Communications support for all coaxial and hybrid fiber-coaxial cable TV network tree and branch topologies. (See Figure 1)

- Support of symmetrical and asymmetrical rates
- Support of *Operation, Administrations, and Maintenance* (OAM) functions
- Support of one-way delays on the order of 400 microseconds (round-trip delays to 800 microseconds)
- Support of a large number of users
- Support for moving data from an originating subnetwork to a destination subnetwork, which may be the same or a different one

Figure 1: CATV Tree and Branch Network

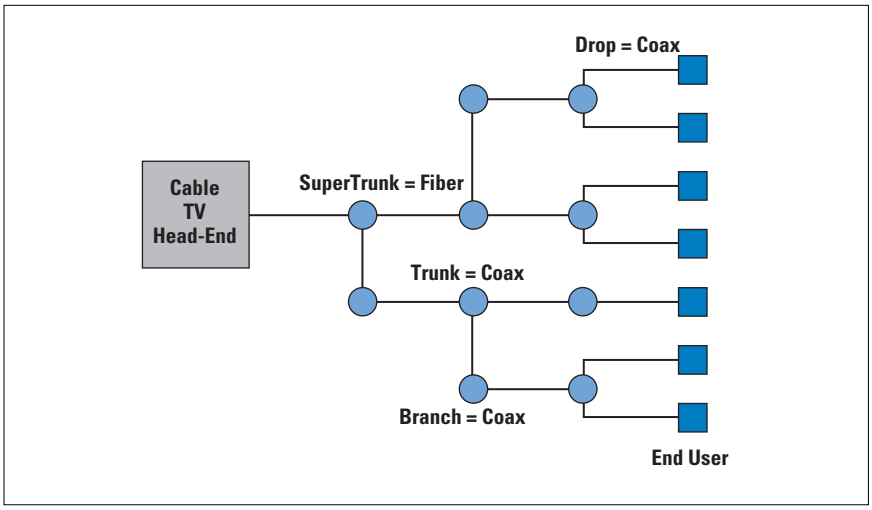

The working group completed a first-release revision of a functional requirements document back in 1995[12], which detailed the 802.14 cable topology model; defined key assumptions, constraints, and parameters; defined key performance metrics and criteria for the selection of multiple PHY protocols and a MAC protocol; and defined the support of *Quality-of-Service* (QoS) parameters. The working group's work plan called for the close of formal proposals in November 1995, with the recommended protocol defined in July 1996. Seventeen MAC protocol proposals were submitted to the working group. Needless to say, it took awhile for the working group to sort through all the issues and opinions. After much consideration, debate, and wrangling of both solutions and personalities, IEEE 802.14 stabilized on a working group draft in September 1998. This working group draft is now being submitted through the IEEE 802 standard approval process.

The 802.14 MAC and PHY specification includes:

- Definition and operational specifications for cable system Head-End Controller and cable modem Stations. (See Figure 2)
- Support of both connectionless and connection-oriented services

- Support of a formal QoS for connections; support for dynamically allocated bandwidth for different types of traffic, including *Constant Bit Rate* (CBR), *Variable Bit Rate* (VBR), and *Available Bit Rate* (ABR)
- Support for unicast, multicast, and broadcast services; interoperability with ATM
- Predictable low-average access delay without sacrificing network throughput
- Fair arbitration for shared access to the network within any level of service
- Downstream channel support for 64 QAM or 256 QAM modulation
- Compatibility for both international and North American downstream digital video standards
- Upstream channel support for QPSK or 16 QAM modulation

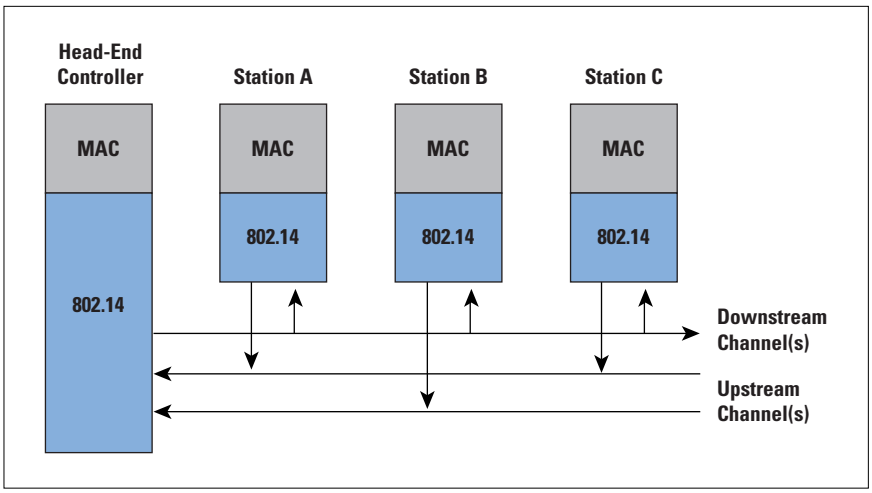

Figure 2: IEEE 802.14 General Model

The selection of ATM cells as the data link layer protocol data unit for IEEE 802.14 networks has the advantage that it provides a suitable integrated multiplexing platform capable of supporting a mix of guaranteed (predictive) traffic flows with best-effort (reactive) traffic flows. See Figure 3. Cable operators can deploy IEEE 802.14 based ATM systems as part of an evolutionary path to a fully integrated multimedia bearer service offering. A residential ATM bearer service easily supports Internet access to the home via the Classical IP over ATM standards of the Internet Engineering Task Force[13] or by providing an IP over Ethernet adaptation overlay service[14]. The development of QoS scheduling support in the Head-End Controller is left for vendors to implement[15, 16, 17].

Figure 3:
IEEE 802.14 ATM
Protocol Model

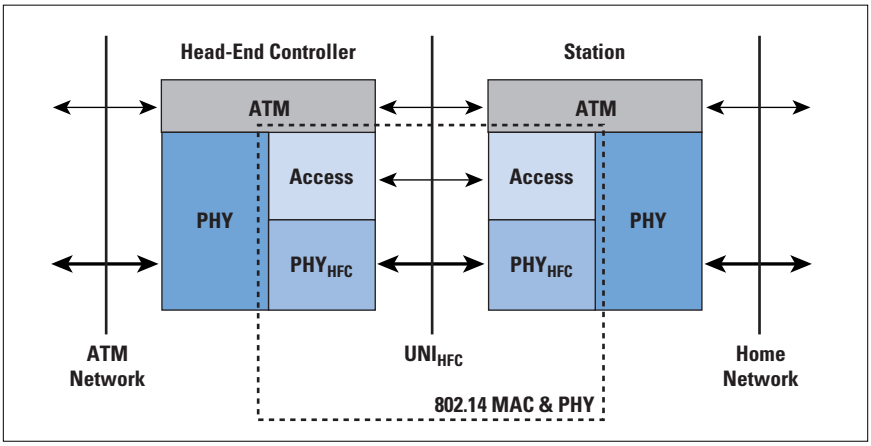

IEEE 802.14 Status

At the time of this writing, the IEEE 802.14 working group just finalized a working group draft suitable to introduction into the IEEE standards process. The entire IEEE process takes about a year from acceptance of the working group letter ballot to producing a published standard.

MCNS DOCSIS

The DOCSIS project is an activity of major cable companies and selected vendors to rapidly develop, on behalf of the North American cable industry, the necessary set of communications and operations support interface specifications for cable modems and associated equipment. The activity was triggered by John Malone in December 1995, in response to competition, vendor postures, and unfortunate lack of progress in the public standards process (that is, IEEE 802.14). The target for the specification was to produce a residential, "low-cost," off-the-shelf, Internet access service, with wide-scale vendor interoperability for base functions with sufficient hooks and room for vendor differentiation.

MCNS specifications are intended to be non-vendor specific, allowing cross-manufacturer compatibility for high-speed data communications services over two-way HFC cable television systems. MCNS met its specification release deadline and published versions of the *DOCSIS Radio Frequency (RF) Interface Specification* V1.0. The first draft specification was published in December 1996. The latest specification was published in July 1998[3]. The DOCSIS RFI protocol is based on the original LANCity symmetric 10 Mbps protocol, evolved to an asymmetric system, with multiple upstream and high-speed downstream (for example, 30 Mbps) channel support.

The MCNS system model is very similar to the IEEE 802.14 general model and includes many interfaces to a cable modem system, as shown in Figure 4. The goal of the DOCSIS project is to produce specifications for the CATV RF interfaces, including behavior of the *Cable Modem Termination System* (CMTS) and *Cable Modem* (CM) with respect to delivery of the residential IP over Ethernet service.

Figure 4: Data-Over-Cable RFI Reference Architecture

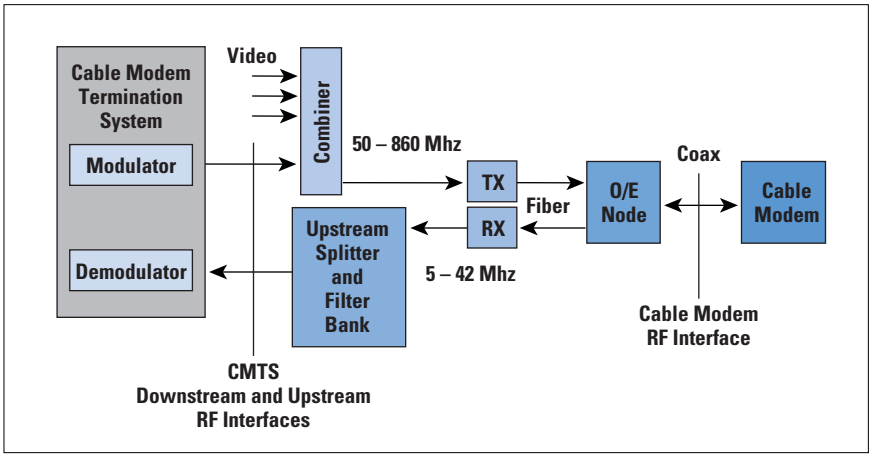

The DOCSIS RFI system is asymmetric, with one to several downstream channels operating asymmetrically with one to several upstream channels. Specific features of MCNS DOCSIS RFI Version 1.0 include:

- Switched Ethernet service for Internet transport via a variable length MAC packet protocol
- Best-effort service
- Downstream data channel rates from 20 Mbps (16 QAM) to 40 Mbps (256 QAM) with a typical configuration of 30 Mbps (64 QAM) in 6 MHz channels
- Compatibility for North American downstream digital video standards. (See article starting on page 27.)
- Downstream data channel rates selected from 320 Kbps (QPSK) through 10.24 Mbps (16 QAM). Channel spectral widths from 200 KHz to 3.2 MHz
- Software flexibility: ability to download new software to change/update CM behavior
- Many filters and features for controlling packet flow and classification
- Comprehensive MIB specifications for control of the cable modem and cable modem termination system
- A single large LAN segment

Due to the time-to-market push for DOCSIS RFI V1.0 interoperable modems, little to no attention was been given for QoS needs however, vendors will likely include some QoS support in their offerings. (Upstream packet fragmentation was removed from the December 1996 draft release.)

CMs and the CMTSs have basically the same protocol stack: downstream and upstream PHY, the DOCSIS RFI MAC, Ethernet and an Ethernet switching layer with substantial filtering, IP/*Address Resolution Protocol* (ARP), *User Datagram Protocol* (UDP), and *Simple Network Management Protocol/Dynamic Host Configuration Protocol/ Trivial File Transfer Protocol* (SNMP/DHCP/TFTP).

The DOCSIS RFI includes upstream and downstream optional packet encryption using the *Data Encryption Standard* (DES) to provide link privacy. RSA public key exchange is used between the CM and CMTS.

DOCSIS RFI Status

CableLabs is actively driving multiple vendor interoperability with the goal of having "silicon interoperability" as soon as possible for DOCSIS "certified" CMs and CMTSs. CableLabs runs a variety of test and certification laboratories in their facility. Numerous vendors are participating. It was the expectation to have many cable modem vendors certified by the cable industry major trade show, the Western Cable Show, in December, 1998. However, as interoperability does take time to work out, the process is taking longer than expected. There will likely be some certified vendors by December 1998, with many more in first quarter 1999. It is now expected that the first widespread deployments of DOCSIS cable modems will start in late first quarter 1999.

The DOCSIS project is currently updating the RFI Version 1.0 specification to include better support for bandwidth management and QoS support. The changes being studied include support for multiple *Service Identifiers* (SIDs), filters to perform the classification of IP packets to different SIDs for differentiated services (QoS), and the signaling support for dynamic SID creations and deletion. A scheme for packet fragmentation will be included which will give substantially better support for managing jitter for delay sensitive traffic, such as packet voice. The primary motivation for adding these extensions to DOCSIS RFI V1.0 is to provide for better support of packet voice and video over DOCSIS IP services. A major focus of the North American cable industry is to support "near toll quality" voice and video services via DOCSIS systems. The cable industry effort writing specification for packet voice and video is called *PacketCable*[18]. It is expected that the DOCSIS RFI V1.1 and initial PacketCable specifications will appear in December 1998.

DOCSIS RFI Version 1.0 was adopted by the *Society of Cable Television Engineers* (SCTE) Data Standards Subcommittee in July 1997 as the North American residential cable modem system standard.

Substantial work is in progress in the IETF *IP over Cable Data Networks* (ipcdn) working group to standardize the DOCSIS MIBs[19, 20] and to standardize IP over DOCSIS[21].

An IP over Cable Modem Example

This section presents a brief overview of a hypothetical IP over HFC system. It is meant to be an informative example to illustrate the application of the IP technology and some of the issues that surround provision of the service over a residential cable TV network. Moving IP datagrams in and out of the home over the cable plant is the important issue. The specific technology and protocols used by the cable modem vendor are important only in their ability to provide required IP service support.

For this example, consider a system that has the following design goals and requirements:

- One-to-many service will be supported in the downstream direction; that is, many cable modems are reachable via the downstream channel
- Many-to-one service will be supported in the upstream direction; that is, the upstream channel bandwidth will be shared. There may be up to several upstream channels
- The protocol used between the Head-End Controller and the head-ends is not significant as long as it meets the needs of the IP service
- The head-end owns the upstream bandwidth and allocates resources to cable modems
- IP over Ethernet 10BaseT is the required interface in the home
- IP over Ethernet or IP over ATM is the required interface at the head-end

This example will rely on the DOCSIS RFI information presented previously in this article. The CMTS can transmit packets to any cable modem on the channel in any order or rate appropriate to the scheduling information it has and controls. The CMTS also participates in the IP multicast group membership (*Internet Group Management Protocol* [IGMP]) and IP *Resource Reservation Protocol* (RVSP) and makes changes in the cable modem resource assignments and allocations as needed. The home cable modem is permitted to use only the upstream channel under direction of the CMTS. Guaranteed and best-effort bandwidth allocations are dynamically assignable by the CMTS. It is assumed that the cable modem protocol has a bandwidth request facility that allows a CM to ask the CMTS for bandwidth. The function of the bandwidth management process is to sort these requests for service and give fair access to the requesting cable modems.

The method for implementation of an Ethernet and 802.3 bridging function over DOCSIS essentially permits the RF channels to act as a serial connection between a half-bridge function in each cable modem with a master in the CMTS. Figure 5 illustrates the protocol stack for this solution. The system presents an Ethernet-like segment to the cable operator. It is well-known how to put together such segments to construct larger internetworks.

Figure 5: Bridged Ethernet via DOCSIS Example

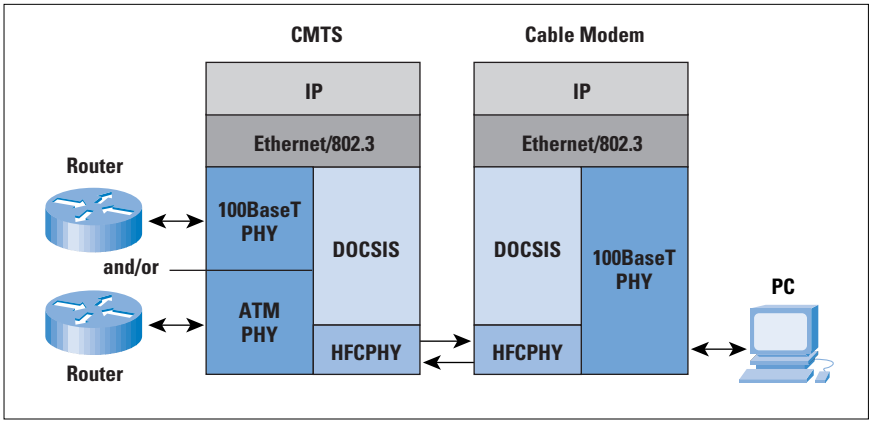

Cable modems provide demarcation between the Internet Service Provider's network and each home network. To help the Internet Service Provider offer fair access service to its residential customers, the cable modem will require sufficient dynamic functionality for multilayer protocol filtering and various forms of rate management (see Figure 6). The goal of this filter is to create a defense perimeter at the first point of entry to the cable network; this perimeter will protect the upstream channel from being saturated or abused by misbehaving home networks. Some examples of this filtering functionality include, but are not limited to:

- Filtering on Ethertype for permitting only certain protocols to pass upstream; for example, IP and ARP only
- Filtering on IP source or destination address to permit/deny access from the home network
- IP and Ethernet broadcast rate limiting; that is, keep any home network broadcast storms confined to the home network
- IP Multicast group address filtering; that is, explicitly permit participation of the home network in an IP multicast group

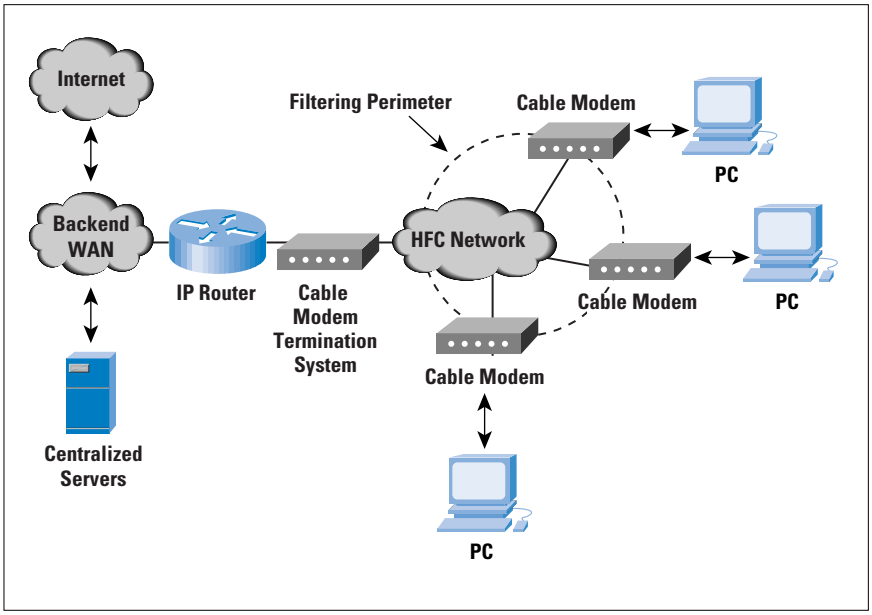

Figure 6: Internet Services via Cable Modem Deployment Model

It should be noted that these filtering functions are under consideration by numerous cable modem manufacturers, and they are being discussed in the IETF ipcdn working group.

A brief overview of IP over cable TV networks has been presented. From an engineering and deployment viewpoint, making the Internet move over cable modems is deceptively straightforward. Many issues are beyond the scope of this article: address allocation methods, back-end network design, configuration services, server placement, home customer support services, installation, firewalls, and troubleshooting.

Summary

This article has presented an overview of the work in progress of the IEEE 802.14 Cable TV MAC and PHY Protocol Standards working group and the MCNS DOCSIS effort. Initial review of these works is positive; indications are that data over HFC systems are viable. The IEEE 802.14 effort began as a study group in late 1993 and has yet to produce a standard. The MCNS DOCSIS process started in early 1996, moved rapidly, and has produced an accepted international standard specification for North American cable operators for residential cable modem service. The IEEE 802.14 standard appears to be destined for some international use and in systems where ATM over CATV is preferred by cable operators.

The cable network environment will provide a very usable and scaleable bandwidth platform for delivering Internet services to and from the home[22]. A hypothetical example was provided that illustrates a general equipment deployment model. Actual deployment of Internet to the home will occur in many areas of North America in 1998 with increasing and substantial deployment in 1999.

For More Information

Information on the IEEE's 802.14 working group can be found on the World Wide Web at: **http://www.walkingdog.com/**

Information the Internet Engineering Task Force's IP over Cable Data Networks working group can be found at: **http://www.ietf.org/**

Information on the North American MCNS DOCSIS effort can be found at: **http://www.cablemodem.com/**

Information on the North American PacketCable effort can be found at: **http://www.packetcable.com/**

Information on the SCTE Data Standards Subcommittee can be found at: **http://www.cablenet.org/scte/scte_dcs.html**

References

[1] Baran, Paul, "On Distributed Communication Networks." IEEE *Transactions on Communication Systems*, Vol. CS-12, pp. 1–9, March 1964.

[2] ATM Forum, "ATM User-Network Interface Signaling 4.0," Specification number af-sig-0061.000, **www.atmforum.com**, July, 1996.

[3] MCNS, "Data-Over-Cable Service Interface Specification—Radio Frequency Interface." SP-RFI-I02-981008, **www.cablemodem.com**, July, 1998.

[4] MCNS, **www.cablemodem.com**, main page, April 1998.

[5] Kim, Albert. "Two-Way Plant Characterization." Technical Session 23, National Cable Television Association Show and Conference, Dallas, Texas, May 9, 1995.

[6] Chelehemal, M., Prodan, R., et al., "Field Evaluation of Reverse-Band Channel Impairments." Society of Cable Telecommunications Engineers, Emerging Technologies Conference, San Francisco, California, January 9–12, 1996.

[7] Laubach, Mark, "Avoiding Gridlock on the Data Infobahn: Port Mismatches Pose Challenges." *CED Magazine,* March 1998

[8] Abramson, Norman, "Development of the ALOHANET." IEEE *Transactions on Information Theory,* Vol. IT-31, pp. 119–123, March 1985.

[9] XEROX, "The Ethernet, A Local Area Network: Data Link Layer and Physical Layer Specification." X3T51/80-50, Xerox Corporation, Stamford, Connecticut, October 1980.

[10] IEEE, "Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications." Standard 802.3-1985 (ISO DIS 8802/3), IEEE, New York, ISBN 0-471-82749-5, 1985.

[11] IEEE, "IEEE Standards for Local Area Networks: Logical Link Control, ANSI/IEEE Std 802.2-1985." Fifth printing, February 1988.

[12] IEEE 802.14 Working Group, "Cable-TV Functional Requirements and Evaluation Criteria." Work in progress, IEEE802.14/94-002R2, IEEE 802 Committee, February 1995.

[13] Laubach, Mark. "Classical IP and ARP over ATM." RFC 1577, January 1994.

[14] Laubach, Mark, "Logical IP Subnetworks over IEEE 802.14 Services." Work in progress, **draft-ietf-ipcdn-ipover-802d14-01.txt**, November 1997.

[15] Laubach, Mark, "Serving Up Quality of Service." *CED Magazine*, April 1997.

[16] Laubach, Mark, "Deploying ATM Residential Broadband Networks." NCTA Cable 96 Conference, Los Angeles, California, April 30, 1996.

[17] Nichols, Kathleen, and Laubach, Mark, "On Quality of Service in an ATM-based HFC Architecture." IEEE ATM Workshop 96, San Francisco, California, August 27, 1996.

[18] PacketCable, "What is PacketCable?" **http://www.packetcable.com**, April 1998.

[19] Roeck, Guenter, "Cable Device Management Information Base for MCNS compliant Cable Modems and Cable Modem Termination Systems." Work in progress, **draft-ietf-ipcdn-cable-device-mib-05.txt**, October 1998.

[20] Roeck, Guenter, "Radio Frequency (RF) Interface Management Information Base for MCNS compliant RF interfaces." Work in progress, **draft-ietf-ipcdn-rf-interface-mib-05.txt**, October 1998.

[21] White, Gerry, "Logical IP Subnetworks over MCNS Data Link Services." Work in progress, **draft-ietf-ipcdn-ip-over-mcns-00.txt**, August 1997.

[22] Lucien Rhodes, "The Race for More Bandwidth." (Interview with Milo Medin of @Home), *Wired Magazine*, Vol. 4.01, January 1996

Internet Drafts are *works in progress* and can be retrieved from:
ftp://ds.internic.net/internet-drafts

MARK LAUBACH holds a B.E.E. and M.Sc. from the University of Delaware. He is Vice President and Chief Technical Officer at Com21, Inc. in Milpitas, California, and is responsible for the end-to-end systems architecture and ATM over HFC protocol specification of the Com21 product family. Prior to Com21, he was with the Hewlett-Packard Company for 14.5 years. Laubach is a member of the IETF, and is past chair of the IP over ATM working group. He is the author of RFC 1577, "Classical IP and ARP over ATM." He regularly attends IETF, IEEE, and SCTE working group meetings. He is a Senior member of the IEEE and a member of the SCTE. E-mail: **laubach@com21.com**

Digital Television: A New Venue for the Internet

by George Abe, Cisco Systems

The digitization of television is of interest to the Internet community in that it opens the possibility of a new mode of delivering IP packets to the home. IP services can be delivered over television broadcast distribution networks, whether over the air, cable, or satellite. This article introduces the basic concepts of *digital television* (DTV) and provides a point of departure for further reading.

Why Is Digital TV Happening?

The original motivation for the research into advanced TV (we avoid the term DTV for a moment) was to prop up sagging TV sales. It was mostly vendor push.

By the late 1970s, Japan and Korea had achieved domination in the production of TV sets worldwide. They were so successful that the market had become saturated, particularly in the developed world. Everyone had one or, more likely, three or four TVs at home. Further, a TV lasts over 10 years, so the replacement market is low. TV production had ceased to be a growth market. Margins were and are poor and few innovations were on the horizon.

So in the early 1980s Japan had begun research into new high-definition televisions that would stimulate new demand and enable them to keep their market leadership. Their system is called *Multiple Subnyquist* (MUSE). MUSE was an analog system, but it had better-quality pictures.

Not to be outdone, the U.S. decided it needed to try to recapture the TV market, so began its own development, under the aegis of the Federal Government. A partnership called the *Grand Alliance* was formed, and it began working in 1984. Pioneering work was done by the partnership members, particularly Zenith, MIT, and General Instruments. They created a digital specification after more than a decade of research and development. Along the way, the computer industry made contributions (or some would say interferences) of its own until the FCC announced a final specification in December 1996. The basic elements are found at **`www.atsc.org`** and referenced later in this article.

Benefits of DTV

The movement toward widespread DTV gained momentum among government officials, broadcasters, and hardware vendors when some of the benefits became clear.

First, because of improvements in technology, it is possible to transmit pictures and sound of significantly higher quality in the same 6 MHz spectrum that analog TV occupies. The 6 MHz spectrum is wasteful of bandwidth, and the government would like to recover the excess so it can be auctioned or used to support other public services (police, fire,

deep space probes, and so on), which could operate at the relatively low frequencies of VHF TV.

Second, digitally encoded TV could provide new services, such as Web access via TV or interactive TV. These have long been dreams of the consumer electronics (CE) industry, but hope springs eternal.

Third, digital TV offers greater security to the programmer and the network. There is a cottage industry in hacking analog set-top boxes. Digital techniques, such as the *Data Encryption Standard* (DES), double DES, and triple DES give operators hope that they can secure their pay-per-view content.

Finally and most interestingly, since digital TV occupies less bandwidth per program, broadcasters, satellite operators, and cable operators have the opportunity to offer more channels. Instead of a mere 10–13 channels available over the air in a single metropolitan area, it is possible to have perhaps 60 or more over the air channels. Cable operators, with their greater bandwidth underground, could have many more channels. Although technically cable could offer 500 channels, it is hard to imagine where the scripts would come from.

What Is DTV?

By our definition, digital television is the capture, production, distribution, and broadcast of programming in a digitally encoded format. Whereas today's analog TV transmits in amplitude modulation, DTV would use *Quadrature Phase Shift Keying* (QPSK), *Quadrature Amplitude Modulation* (QAM), or *Vestigal Side Band* (VSB) modulation techniques. We won't detail these techniques here except to mention that they are mutually incompatible.

When DTV standards were discussed in the 1980s, the industry could not agree on a single display. The deliberations became more protracted with the entry of the computer industry into the discussions, long after the broadcasters and consumer electronics people began their work. Would there be interlaced or progressive scanning? Would there be the existing aspect ratio or would there be a wide-screen display? Square pixels or not? How many lines of resolution would be displayed?

With the broadcasters and consumer electronics vendors arguing for interlacing, oval pixels, and wide screens and the computer people arguing for progressive scanning, square pixels, and a more square display, the disagreements could not be bridged.

Therefore, the FCC had no choice but to declare that the "market should decide" which display format would prevail. Accordingly, the FCC announced in December 1996 that 18 different display formats would be permissible for over-the-air digital TV. A broadcaster could elect to transmit in any of the approved formats. The approved formats are shown in Tables 1 and 2.

Table 1: Progressive Video Scanning Formats for Digital TV

Vertical Lines	Horizontal Pixels	Aspect Ratio	Frame Rate per Second
1080	1920	16:9	24, 30
720	1280	16:9	24, 30, 60
480	704	16:9	24, 30, 60
480	704	4:3	24, 30, 60
480	640	4:3	24, 30, 60

Table 2: Interlaced Video Scanning Formats for Digital TV

Vertical Lines	Horizontal Pixels	Aspect Ratio	Frame Rate per Second
1080	1920	16:9	30
480	704	16:9	30
480	704	4:3	30
480	640	4:3	30

The vernacular to describe the formats typically indicates the number of vertical lines and the scanning format. For example, "1080i" refers to 1080 lines, interlaced scanning; "720p" refers to 720 lines in progressive format.

In practice, only a few of the 18 approved formats are under consideration by the nation's broadcasters. NBC and CBS have declared they will support 1080i. ABC is opting for 720p, and Fox has opted for 480p.

Apart from the controversy over display, most of the other elements were quickly resolved. Modulation scheme, transport multiplexing, compression, timing, and an overall systems and testing procedure were agreed to. The apparatus for DTV was in place, almost. The time was January 1997.

High Definition or Standard Definition

Some view DTV as synonymous with high-definition television. It is not. DTV encompasses both *High-Definition TV* (HDTV) and *Standard-Definition TV* (SDTV). Hence HDTV is a proper subset of DTV. The difference between HD and SDTV is not standardized, but our definition of HD includes the display formats that have 720 or 1080 lines. Formats with fewer lines are standard definition.

The key point of difference between HD and SD is that with HD and current compression techniques (MPEG-2), only one program is accommodated in one 6-MHz channel. With SD, it is possible for the broadcaster to transmit two or more programs simultaneously, in a single 6-MHz chunk of bandwidth.

This has tremendous implications. If broadcasters can transmit multiple channels at once, it would be possible (technically) for Disney to broadcast ABC, the Disney Channel, ESPN, and A&E over the air in the same bandwidth they use to show ABC today. (Of course they won't do this for commercial and contractual reasons, but the technology makes it doable).

For Internet Service Providers, a broadcast could transmit SD programming simultaneously with datacasting, and go into the push-mode data service business. For example, Disney/ABC could download software updates for Disney Interactive, or perhaps contract with Microsoft to deliver Windows updates. Whereas most Internet folk view MPEG being transported inside IP packets on the Internet, broadcasters intend to insert IP packets into MPEG-2 transport streams. The consumer's digital set-top box would tune to the data "channel," extract the data from its MPEG capsule, and divert the data packet to an Ethernet or ATM port on the set-top.

There are nearly 1,600 broadcasters in the U.S. Each could, in theory, transmit 19.3 megabits per second. Of course, most of these bits will be used for television, but certainly 1 or 2 megabits can be accommodated by each broadcaster for data service.

Given the dearth of programming to fill multiple SD channels, broadcasters are strongly motivated to consider data services and compete for a slice of the Internet service market.

Digital TV—End to End

Whereas one easily thinks of DTV as a distribution and display technology, in fact there are major changes required to capture, edit, and distribute digital content. Thus there is the need for new cameras, post-production editors, sound mixers, and the like.

Digital TV can be transmitted over the air, through cable networks, or via *Direct Broadcast Satellite* (DBS). Today, only DBS has achieved large-scale distribution of digital TV, with over 7 million subscribers in the U.S. and 15 million worldwide.

Content is created either through a digital camera or by converting existing analog content, such as 35mm film, into digital format. Within the production environment, editing changes are made, typically using *Nonlinear Editors* (NLEs) that connect to a local-area network.

Original production is normally done in the high definition. The highest form of resolution is 1.492 Gbps. (See Table 3.) Equipment to do this is not widely available, but it will be eventually. Panasonic is shipping a digital camera capable of 1.5-Gbps output, but rumor has it they cost almost $500,000, if you can even get one. Nonetheless, 41 stations began HD programming in November, highlighted by an NFL game on CBS between the Buffalo Bills and the New York Jets on November 8.

Some compression is applied within the postproduction and editing environment. The TV industry, through the *Society of Motion Picture and TV Engineers* (**www.smpte.org**), developed a series of digital transmission standards. Chief among these is SMPTE 305M, which defines a protocol called *Serial Data Transport Interface* (SDTI), which calls for a 270- or 360-Mbps service to link various pieces of production equipment such as NLEs in a postproduction facility. SMPTE 305M is a networking scheme complete with an addressing specification.

(Interesting point about 305M: It is the first and only protocol known to this author that specifies use of IPv6 addressing.)

Another important protocol is SMPTE 259M, which is a link-layer protocol underneath 305M.

A competing protocol to SDTI is the *Digital Video Broadcasters Asynchronous Serial Interface* (DVB-ASI). Information on DVB-ASI is found at **www.dvb.org**.

From the editing environment, content is distributed via satellite or land lines to local affiliates (for local over-the-air broadcast), cable head-ends (for cable TV distribution) and satellite hubs (for direct-to-home satellite service). The distribution from national feeds to local facilities is normally at T3/E3 speeds because of the availability of T3/E3 services by telephone companies and satellite transponders for affiliate and direct-to-home distribution.

Cable providers, local broadcasters, and satellite services add their own content and make certain changes to the national feeds. Among these changes are assignment of the programming to specific frequencies or channels, insertion of local advertising, local programming, and emergency broadcasts.

After adding their own content, the local services distribute the final programming to consumers. Over-the-air broadcasters will transmit 19.3 Mbps per 6 MHz, cable will transmit 27 Mbps per 6 MHz, and satellite uses variable channelization, kept closely under wraps.

So there is the progression downward from 1492 Mbps of original encoding, to 270 Mbps for editing, to 34/45 Mbps for affiliate distribution, to 27 Mbps or less for distribution to the end user.

Table 3: Bit Rate Requirements for Various Display Formats

Format	Pixels per Line	Lines per Frame	Pixels per Frame	Frames per Second	Millions of Pixels per Second	Bits per Pixel	Mbps
SVGA	800	600	480,000	72	34.6	8	276.5
NTSC	640	480	307,200	30	9.2	24	221.2
PAL	580	575	333,500	50	16.7	24	400.2
SECAM	580	575	333,500	50	16.7	24	400.2
HDTV	1920	1080	2,073,600	30	62.2	24	1492.8
Film	2000	1700	3,400,000	24	81.6	32	2611.2

Note: Film display formats vary, depending on content and directorial prerogative.

Over the Air and Cable

All the huffing and puffing by the FCC, the consumer electronics industry, the computer industry, and the broadcasters pertains to over-the-air transmission. However, about two-thirds of the American viewing public views TV through cable. So if most Americans are to receive DTV, they must receive it through cable.

This raises important technical and regulatory questions. The technical question is: How are the digital signals produced by the broadcasters and their affiliates to be sent through wires, and what is the allocation of functions between the digital set-top and the digital receiver? This question seems simple but it is not, as we shall see.

The regulatory question pertains to whether the cable operators are to be compelled to carry DTV from broadcasters. This problem is referred to as the digital *Must Carry Problem,* now under consideration by the FCC. It certainly will be litigated, whatever the outcome of the FCC's decision.

Technical Question

Among the key provisions agreed to by the Grand Alliance is the use of a modulation technique called 8-VSB for over-the-air digital transmission. The particulars of 8-VSB are not significant here, but we will mention that this particular decision was arrived at in the mid-1980s, before the cable industry had much impact on the viewing public or on the broadcasting industry.

When the cable industry began to think about digital, in the mid-1990s, they settled on a modulation scheme called 64 QAM. 64 QAM is able to produce 27 Mbps in 6 MHz, whereas 8-VSB produces about 19.3 Mbps. The difference occurs because over-the-air broadcasting requires a more robust encoding scheme to combat the more hostile nature of over-the-air transmission, as opposed to the safer environment of coaxial cables. Thus the cable modulation technique can be more aggressive than over-the-air techniques.

(We should add that satellites use an even more robust modulation technique called QPSK, which gets fewer bits per Hertz than VSB or QAM. But robustness is needed because satellite signals must travel far greater distances than cable or local broadcast.)

Thus for cable to carry a digital over-the-air broadcast, some conversion of 8-VSB encoding to 64 QAM encoding is necessary. This necessity does not present a major technical problem, but agreement is needed on where the conversion is done and at what cost. For example, Broadcom and Sony are collaborating on the development of a chip, to be embedded in a TV, that can decode VSB and QAM. It sounds simple, but the cable industry is not interested. They want to carry QAM and QAM only on their networks.

One option is to convert the format of the digital bitstream coming out of the cable box to the IEEE 1394 *FireWire* format. Since DTVs are likely to have FireWire input, this conversion can provide a ubiquitous connection. However, this scenario raises the problem of copy protection, a sore point in Hollywood. Since digital copies are pristine, the content providers (studios and record companies) are firm in their resolve that unless there is strong copy protection, none of their content will be available over FireWire.

Another option is to build a set-top box that takes baseband signals and modulates them to look like 8-VSB broadcast signals on channel 3, similar to how VCRs work in the analog world now. This scenario is clearly rather ugly, but understood by consumers.

Finally, it could be up to the cable operators to transmodulate the 8-VSB into QAM at the cable head-end. Better yet, they can accept broadcasters' feeds in baseband, and then QAM-modulate the baseband signals for their consumers. The cable set-top box would be sending bit maps to a dumb digital monitor, like a computer monitor, which doesn't know or care that it is receiving QAM or VSB programming.

Apart from modulation, there is the issue of display format. NBC and CBS have declared they will transmit in 1080i. ABC has chosen 720p and Fox has chosen 480p, with some vague pledge for higher definition later. After all, it does not seem necessary to show *The Simpsons* in HD.

On the other hand, John Malone, Chairman of TCI, went public in May 1998 with his declaration that TCI would not voluntarily carry 1080i because it (1080i) was wasteful of bandwidth. Implied in his comment is the fact that cable operators do need to be restricted to 6-MHz channelization for digital. In fact, the entire DTV spectrum on cable could be considered a gigantic pool of bandwidth that the cable operator could allocate to individual channels, much as direct satellite does. This setup gives the cable operators incentive to downconvert the broadcasters' DTV signals. For example, when NBC sends 1080i, the cable operator may elect to transmit 720p, or less, to its customers.

Should the cable operators be required to carry the HDTV pictures from the broadcasters in the broadcasters' chosen format? Would they be allowed to downconvert the HD into standard definition? What happens when a broadcaster, say NBC, elects to transmit in SDTV and thereby has the capability of multiplexing several channels onto a single chunk of 6 MHz? What is the duty of the cable operator to carry Internet datacasting offered by the broadcasters over the cable network, in competition with services such as @Home and Roadrunner?

The complexities of multiplexing go further. Let's say ABC elects to broadcast SD. If one of the subprograms in the multiplex is a pay-per-view channel, should the authentication procedures of the cable operator be superceded? Should the electronic program guide of the cable operator be superceded?

Questions like these have technical and regulatory aspects and are being worked in industry, the FCC, and state regulatory agencies. It is possible that Congress will get involved as well. When John Malone made his statement, both sides of the aisle in Congress were not amused. They want DTV to happen so that spectrum can be freed. If the cable operators stand in the way, the conversion to digital is stopped dead in its tracks.

The Open Cable Initiative

The cable industry does not want to be a bottleneck to broadcasters. On the other hand, it needs to make quick progress into DTV to compete against satellite. Therefore, the industry has embarked on a process called *Open Cable,* which seeks to define a digital set-top box that can be available at retail. Available at retail means a nonproprietary, open design. Open Cable strives to make the DTV set-top box independent of processor platform (that is, not an Intel Pentium necessarily) and operating system independent (that is, not a Microsoft Windows CE necessarily).

The Open Cable set-top box will allow for data services through a specification written by the *Digital Audio Visual Council* (DAVIC—**www.davic.org**) and therefore, is not compatible with the current *Data-over-Cable Service Interface* (DOCSIS) specification supported by the U.S. cable industry. (See article starting on page 13.) However, it is possible for DOCSIS capabilities to be added on to an Open Cable set-top box. We mention Open Cable because it will be the key customer premises device for cable and digital TV and much hinges on its interoperability with broadcasters transmissions.

Digital TV via Satellite

In addition to over-the-air and cable, DTV can be received by satellite. As of this writing, it is the only way to receive DTV. The digital satellite industry has nearly 7 million subscribers who received DTV today. Its role in all the discussions of HD vs. SD and the provision of data services is relatively low key because it is believed that satellite will continue to be a niche provider because of its technical and legal problems in distributing locally originated TV stations.

But satellites bear watching because if they are able to deliver local channels and obtain 15–20 million homes in the U.S., then the financial consequences on cable and over the air could be crucial.

The New Digital Studio

The figure shows a schematic of the elements of a DTV broadcast studio described recently by the U.S. *National Institute of Standard and Technology* (NIST). At the heart of the studio is an ATM switch with new interfaces that connect to DVB or ATSC infrastructures via DVB-ASI or SDTI interfaces.

Connection for wide-area distribution will likely be over ATM. Converters exist for DVB-ASI to ATM. For example, Cellware (**www.cellware.de**) in Germany markets such a converter, but there is no SDTI-to-ATM interface known to this author at this time.

The digital studio provides a new a marketing opportunity for the LAN industry. Broadcast digital production demands higher speeds than most other LAN applications.

Thus vendors of data communications equipment have two opportunities: to provide equipment to broadcasters who want to enter the Internet service business and to production houses that use ATM or other LANs to support editing and production applications.

Figure 1: Prototype of HDTV Broadcast Studio

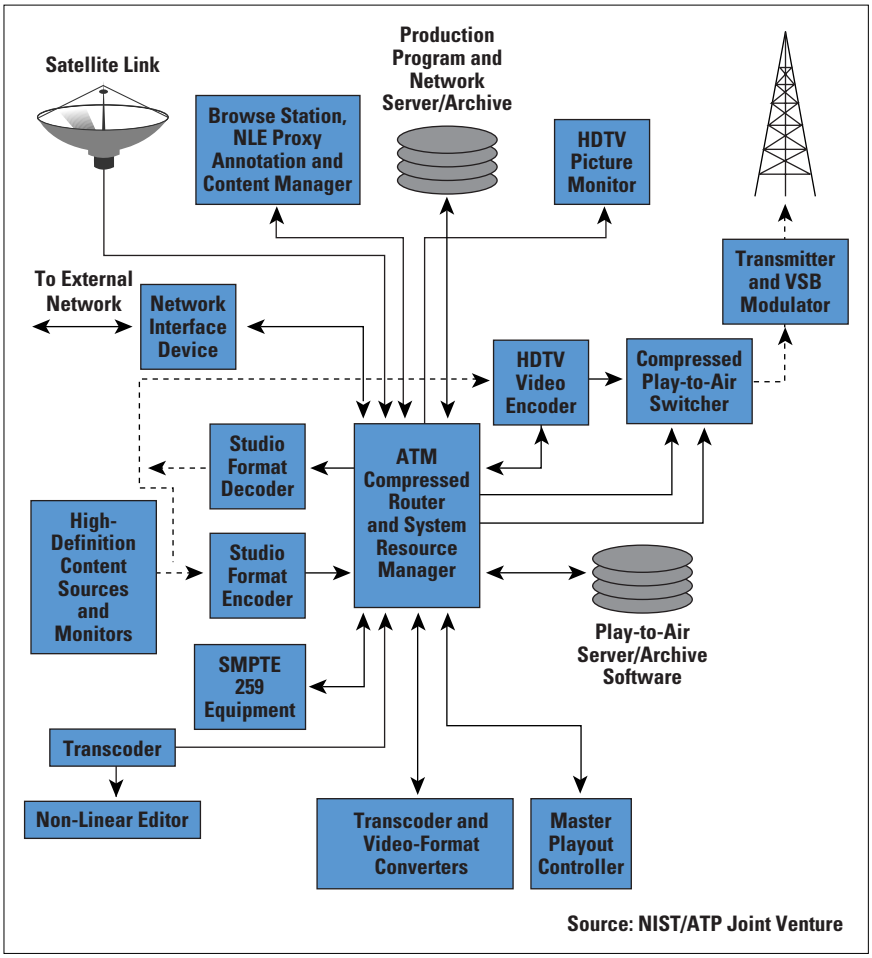

Web Sites

www.atsc.org: *Advanced TV Standards Committee*. S13 and S16 are subgroups working on datacasting; S13 focuses primarily on the downstream path, whereas S16 focuses primarily on the reverse communication from the receiver. Since over-the-air is one way, this work is limited to the communications between the S13 forward channels and a telephone or Internet return path.

www.dvb.org: The *Digital Video Broadcasting Project* (DVB) has taken the lead in defining DTV specifications as well as defining datacasting interfaces over DTV infrastructures.

www.smpte.org: *Society of Motion Picture and Television Engineers*.

www.sbe.org: *Society of Broadcast Engineers*.

www.scte.org: *Society of Cable TV Engineers*.

www.mpeg.org: *Motion Picture Experts Group*. The word on MPEG compression, controls, and transmission.

References

[1] ISO/IEC IS 13818-1, International Standard, MPEG-2 Systems.

[2] ISO/IEC IS 13818-2, International Standard, MPEG-2 Video.

[3] ISO/IEC 13818-6, International Standard, Digital Storage Media Command and Control (DSM-CC).

[4] ATSC Standard A/52 (1995), Digital Audio Compression (AC-3).

[5] ATSC Standard A/53 (1995), ATSC Digital Television Standard.

[6] ATSC Standard A/55 (1996), Program Guide for Digital Television.

[7] ATSC Standard A/56 (1996), System Information for Digital Television.

[8] ATSC Standard A/57 (1996), Program/Episode/Version Identification.

[9] ATSC Standard A/63 (1997), Standard for coding 25/50-Hz Video.

[10] ATSC Standard A/64 (1997), Transmission Measurement and Compliance For Digital Television.

[11] ATSC Standard A/65 (1998), Program and System Information Protocol for Terrestrial Broadcast and Cable.

[12] ATSC T3/S13 Doc. 010 DVS-yyy Rev z Draft, ATSC Data Broadcast Specification for Terrestrial Broadcast and Cable.

[13] ETR XXX: Digital Video Broadcasting (DVB); Guidelines for the Use of the DVB Specification: Network Independent Protocols for Interactive Services (ETS 300 802).

[14] SCTE DVS-nn: SCTE Digital Video Subcommittee (DVS) standard for Cable Headend and Distribution Systems (spec not released—under development)

GEORGE ABE holds an A.B. in Mathematics and an M.S. in Operations Research from UCLA. He currently is a Consulting Engineer at Cisco Systems, where he has dabbled in various areas of residential broadband networking since 1994. He is the author of *Residential Broadband*, Cisco Press (imprint of Macmillan Press). He expects to be a early adopter of digital TV and, when not watching TV, he can be reached at **georgea@acm.org**

I Remember IANA

by Vint Cerf, MCI WorldCom
October 17, 1998

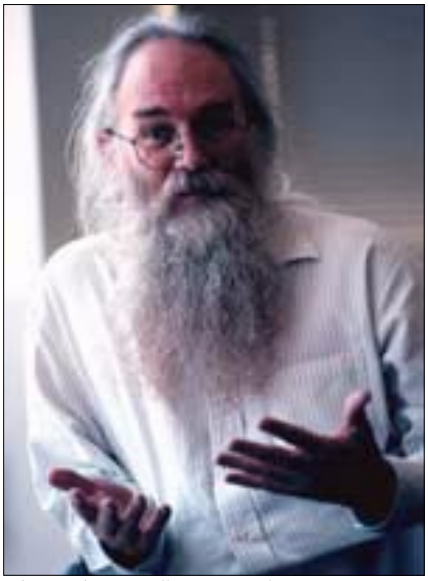

Photo: Chris Pizzello, New York Times Pictures

A long time ago, in a network, far far away, a great adventure took place! Out of the chaos of new ideas for communication, the experiments, the tentative designs, and crucible of testing, there emerged a cornucopia of networks. Beginning with the ARPANET, an endless stream of networks evolved, and ultimately were interlinked to become the Internet. Someone had to keep track of all the protocols, the identifiers, networks and addresses and ultimately the names of all the things in the networked universe. And someone had to keep track of all the information that erupted with volcanic force from the intensity of the debates and discussions and endless invention that has continued unabated for 30 years. That someone was Jonathan B. Postel, our *Internet Assigned Numbers Authority* (IANA), friend, engineer, confidant, leader, icon, and now, first of the giants to depart from our midst.

Jon, our beloved IANA, is gone. Even as I write these words I cannot quite grasp this stark fact. We had almost lost him once before in 1991. Surely we knew he was at risk as are we all. But he had been our rock, the foundation on which our every Web search and e-mail was built, always there to mediate the random dispute, to remind us when our documentation did not do justice to its subject, to make difficult decisions with apparent ease, and to consult when careful consideration was needed. We will survive our loss and we will remember. He has left a monumental legacy for all Internauts to contemplate. Steadfast service for decades, moving when others seemed paralyzed, always finding the right course in a complex minefield of technical and sometimes political obstacles.

Jon and I went to the same high school, Van Nuys High, in the San Fernando Valley north of Los Angeles. But we were in different classes and I really didn't know him then. Our real meeting came at UCLA when we became a part of a group of graduate students working for Professor Leonard Kleinrock on the ARPANET project. Steve Crocker was another of the Van Nuys crowd who was part of the team and led the development of the first host-to-host protocols for the ARPANET. When Steve invented the idea of the *Request for Comments* (RFC) series, Jon became the instant editor. When we needed to keep track of all the hosts and protocol identifiers, Jon volunteered to be the Numbers Czar and later the IANA once the Internet was in place. Jon was a founding member of the *Internet Architecture Board* (IAB) and served continuously from its founding to the present. He was the *first* individual member of the Internet Society—I know, because he and Steve Wolff raced to see who could fill out the application forms and make payment first and Jon won. He served as a trustee of the Internet Society.

He was the custodian of the .us domain, a founder of the Los Nettos Internet service, and, by the way, managed the networking research division of USC Information Sciences Institute.

Jon loved the outdoors. I know he used to enjoy backpacking in the high Sierras around Yosemite. Bearded and sandaled, Jon was our resident hippie-patriarch at UCLA. He was a private person but fully capable of engaging photon torpedoes and going to battle stations in a good engineering argument. And he could be stubborn beyond all expectation. He could have outwaited the Sphinx in a staring contest, I think.

Jon inspired loyalty and steadfast devotion among his friends and his colleagues. For me, he personified the words "selfless service." For nearly 30 years, Jon has served us all, taken little in return, indeed sometimes receiving abuse when he should have received our deepest appreciation. It was particularly gratifying at the last Internet Society meeting in Geneva to see Jon receive the Silver Medal of the International Telecommunications Union. It is an award generally reserved for Heads of State, but I can think of no one more deserving of global recognition for his contributions.

While it seems almost impossible to avoid feeling an enormous sense of loss, as if a yawning gap in our networked universe had opened up and swallowed our friend, I must tell you that I am comforted as I contemplate what Jon has wrought. He leaves a legacy of edited documents that tell our collective Internet story, including not only the technical but also the poetic and whimsical as well. He completed the incorporation of a successor to his service as IANA and leaves a lasting legacy of service to the community in that role. His memory is rich and vibrant and will not fade from our collective consciousness. "What would Jon have done?" we will think, as we wrestle in the days ahead with the problems Jon kept so well tamed for so many years.

There will almost surely be many memorials to Jon's monumental service to the Internet Community. As current chairman of the Internet Society, I pledge to establish an award in Jon's name to recognize long-standing service to the community, the *Jonathan B. Postel Service Award*, which will be awarded to Jon posthumously as its first recipient.

If Jon were here, I am sure he would urge us not to mourn his passing but to celebrate his life and his contributions. He would remind us that there is still much work to be done and that we now have the responsibility and the opportunity to do our part. I doubt that anyone could possibly duplicate his record, but it stands as a measure of one man's astonishing contribution to a community he knew and loved.

VINTON G. CERF is senior vice president of Internet Architecture and Technology for MCI WorldCom. Widely known as a "Father of the Internet," he is the co-designer of the TCP/IP protocol. Cerf served as founding president of the Internet Society from 1992–1995 and is currently chairman of the Board. Cerf holds a Bachelor of Science degree in Mathematics from Stanford University and Master of Science and Ph.D. degrees in Computer Science from UCLA. E-mail: **vcerf@mci.net**

Book Reviews

Internet Messaging

Internet Messaging: From the Desktop to the Enterprise, by Marshall T. Rose and David Strom ISBN 0-13-978610-4, Prentice-Hall PTR, 1998, **http://www.prenhall.com**

Very few Internet voices hold a status equivalent to E.F. Hutton's advertising campaign: "When they speak, we should listen." Marshall Rose and David Strom are two such voices, making any product of their combined efforts a serious matter, indeed. Rose has typically written about basic technology, Strom about the pragmatics of use, especially trials and tribulations of fitting networked pieces together. *Internet Messaging* is in the latter category, with a strong added introduction of e-mail and security technology. Anyone who has professional contact with e-mail should get a copy of this book. If commercial use of Internet mail were more advanced and stable, we probably would not need an effort like this. However, e-mail professionals must constantly deal with problems in using interesting functions and in troubleshooting interoperability. *Internet Messaging* helps with the planning, use and debugging of complex, or otherwise "interesting," e-mail services.

Updated Information

The book provides a superb survey of the relevant technology, the popular user mail software, and the rather interesting range of mail and messaging operations issues, including styles of use by organizations. The comparisons of different mail systems leave the reader with a solid understanding of functional and usage requirements for modern systems, as well as the choices available at the time of publication. Mary Houten-Kemp's Web site at **http://www.everythingemail.net** is being used to provide updated information.

E-mail includes a wide range of technical and operations issues, and *Internet Messaging* touches all of them. Its introductions cover user environment, mail transfer, mailing list services, unsolicited bulk e-mail ("spam"), encryption-based security, remote user access, virtual private networks, and directory services. Providing a single discussion, which integrates the use of these disparate technologies, is enough to justify the book.

Organization

Internet Messaging attempts very regular organization and states that the goal is to permit use as a problem/solution reference work. It primarily distinguishes between sending and receiving functions and between desktop and enterprise requirements. This creates a two-by-two matrix, defining the core four chapters. The other chapters include philosophical opening and closing discussions, a separate, very informative chapter on security, and another on general enterprise operations issues.

Most of the chapters are organized into Introduction, Problems, Standards, and Solutions. Unfortunately that regularization is all that is shown in the Table of Contents, so the reader gets little help finding specifics by reading the Table. Similarly, the organization of the chapter contents did not seem compelling for use in problem solving. The additional "How Can I" matrix (on page 10) and its associated discussion text is intended as the primary means for locating relevant discussions.

Comparisons

User software comparisons are given throughout the book, for Microsoft Outlook 4.01, Netscape Messenger 4.04, Qualcomm Eudora Pro 4.0, Lotus cc:Mail 8.1, CompuServe WinCIM 3.02, and America Online 3.0. Specific mailing lists, security, remote access, and directory software and services are also reviewed. Oddly, the discussion of remote access mentions only global, single-provider services—and their favorite is currently having financial problems—but did not mention the "association" style of service that integrates many independent providers, notably GRIC and iPass. (Full disclosure: iPass is a client.)

Most products are undergoing aggressive enhancement so that no printed text can be entirely up-to-date. Hence the Web site. For the software and services I know well, the book looked reasonable. Of course it is not entirely error free, but the errors are small and perfect detail is not required. I believe there are two major benefits to these comparisons. One is that the reader is given a very solid sense of the general capabilities and limitations of modern e-mail software. The second is to make a reasonable, first-pass filtering of candidate packages to be used in an organization. It would *not* be appropriate to attempt selecting among these packages according to subtle differences reported in the book.

Benefits

As one would expect of these authors, a very large, long-term benefit of their efforts is in their many excellent criticisms and suggestions. Unfortunately, many of them are in notes located at the end of each chapter. It's hard to imagine a less-convenient place to put them, since I found myself constantly shifting back and forth between the main text and the notes. It would not have been so irritating if the comments were less interesting; they should have been true footnotes, with easy access on each page. The stellar example of direct utility from these comments is Figure 2.1 on page 38. It shows a systems structure for user software processing of incoming mail. Every vendor should study this discussion carefully and implement it immediately. Please!

—Dave Crocker
Brandenburg Consulting
dcrocker@brandenburg.com

Web Security

Web Security: A Step-by-Step Reference Guide, by Lincoln D. Stein, ISBN 0-201-63489-9, Addison-Wesley, December 1997,
http://www.awl.com/cseng/titles/0-201-63489-9

Whenever the topic of the World Wide Web comes up, you can be sure that some mention of "security" will soon follow. Web users, Web creators, and even Web technology developers are all keenly aware of the security concerns. But what do we mean by "security?" The safety to use a credit card? Keeping a Web site safe from break-ins? Keeping the kids away from online erotica? And whose security are we concerned with, the user's or the Web site operator's?

This book covers most of what we might expect to find under the umbrella of security. In addition to dealing with the broad scope of Web security, the author also tries to cover the topic with sufficient simplicity for the novice and enough detail for the engineer. The good news is that this book succeeds in delivering a single volume that covers all we could possibly expect on the topic, and at levels suited for a broad audience range.

Organization

The author begins by making the distinction between security for the browser, the Web site, and the network between them. This division of the topic forms the basis for the organization of the book. Moving through each of the three parts, the author proceeds from the simple to the complex in a logical, additive order. He discusses topics introduced early in the book from a functional standpoint—how they affect the user. He may cover the same technology in later chapters, but in greater depth, detailing server and network configuration and discussing the underlying technology.

In the first part of the book, the author covers document confidentiality, including standard "text" documents as well as electronic commerce. A major theme in this section is cryptography. The author presents symmetric and public key encryption technologies from a functional standpoint. He presents various encryption standards, with a discussion of their strengths and weaknesses. In another chapter he provides a good primer on the *Secure Electronic Transaction* (SET) protocol handling, as well as other options (*Common Gateway Interface* [CGI] scripts and *Secure Sockets Layer* [SSL]) for credit card order processing.

In Part 2 we are introduced to issues of client-side security. The author devotes a full chapter to an in-depth explanation of SSL services. He also looks at issues associated with active content, and presents technologies such as Java, ActiveX, and other options, along with notes on their respective security implications. Finally, he covers issues of privacy—in this case, the personal privacy of the user. Throughout these chapters, the author emphasizes user-controllable settings such as browser configuration options.

Whereas the author focuses on user involvement in the first two parts, with an appropriate level of technical content, in part 3, targeted to Web masters and system administrators, he introduces the engineering side with an in-depth coverage of server-side security. He covers the two prominent Web-serving operating systems: UNIX and Windows NT, with good attention to various versions of each. Topics include basic system security, access control, and activity monitoring. Other chapters include an excellent discussion of encryption and certificate technology, safe CGI scripting, remote authoring of Web data, and firewalls.

Presentation and Style

The author illustrates his points with good examples. He also presents appropriate sidebar discussions and illustrations, which not only clarify the information, but also provide interest and variety in what could be a very dry volume. Each chapter ends with a listing of resources, both print and "online." Where appropriate, the author includes checklists to help the reader apply the material just covered.

As a result of the practical, well-grounded presentation of material, we are continually able to see practical applicability to our own situation. For example, the author presents us with information about dangers to our privacy, and why that might be important to us. This is immediately followed by clear instruction on changing privacy-affecting settings in various versions of both Netscape and Internet Explorer. The author uses this technique throughout the book, and it is as useful with password management, CGI scripting, or firewall configuration as it is with privacy.

Recommended

Although experts in encryption and other specific security-related technologies will find this book too simple for their personal area of expertise, the strength of the book is not in its coverage of any one area, but in its well-integrated and cohesive coverage of a broad range of interrelated topics. The ability for any reader, first-time surfer or Web guru, to find practical, easily applied information makes this book a required item on any webmaster's bookshelf, and a must-read for anyone who spends any serious time on the Web.

—Richard Perlman
Berkeley Internet Group
perl@berkinet.com

Internet Cryptography

Internet Cryptography, by Richard E. Smith, ISBN 0-201-92480-3, Addison-Wesley, 1998, **www.awl.com/cseng/titles/0-201-92480-3**

The 1990s might easily be known as the decade of the Internet. The Internet came into the mainstream during this decade, a global frontier with frontier problems and rules. Seemingly overnight, everyone from government agencies to Chinese restaurants had a Web presence. Young children exchanged e-mail with their grandparents and friends, a big change from just a few years ago when it was the domain of technologies and a place where everybody knew your name.

The 1990s could also be known as the decade when cryptography became mainstream. Perhaps because of the change in the Internet community, people became more aware of the need to protect the privacy of internetwork communications. Certainly, the U.S. government's attempt to push government control of cryptographic keys in the Clipper controversy helped to move cryptography and its related issues from science journals to the front pages of our newspapers. Today, while not mainstream, terms such as *Virtual Private Networks* (VPNs), *Secure Sockets Layer* (SSL), *IP Security* (IPSec), *Pretty Good Privacy* (PGP), *Secure Multipurpose Internet Mail Extensions* (S/MIME), and related technologies are known among IT professionals, and cryptography is no longer a tool used only by spies and military communication officers.

The Author

Richard E. Smith is well-known to members of various security-related forums on the Internet, as well as to security conference attendees. A security consultant with Secure Computing Corporation, Smith's background is in military-grade security. His experience on the lecture circuit, explaining issues of firewalls, cryptography, and other computer and network security topics, has directly contributed to production of a book on a lofty subject that is reachable by the nonscientist.

Organization

The chapters of this book fall into three groupings: an introduction to the basics of cryptography, its terms, methods, and mechanisms; network encryption and a discussion of VPNs, focusing on IPSec; and finally public key cryptography as it is used with message and file encryption and "Web" transactions.

The discussion in the opening chapter on basics may scare some off; Smith tends to oscillate between various levels of complexity. Consequently, some members of the intended audience of (quoting from the Preface) "people who know very little about cryptography but need to make technical decisions about cryptographic security," may, for example, zone out during the discussion of IP protocols. My suggestion would be to press on, and not worry about the random item that might go over your head. Everything there has a purpose, and the important information will fall into place by the end of each chapter.

If this book ended with Chapter 4, it would still be a useful book. The complex basics of cryptography and the issues that should be of concern to an information security officer are clearly presented and explained. The only area that is given less than adequate coverage is that of key recovery. Smith makes no mention of legitimate business reasons for the recovery of encrypted data if the originator is unavailable (the proverbial question, "What if you got hit by a truck?"), nor does he mention any mechanism other than the escrow of secret keys, although there are other, safer, methods. Of particular use are Smith's explanations of the various cryptographic algorithms and his discussions of safe key lengths and risks.

In the sections on VPNs and IPSec, Smith covers everything from mobile users and remote access, to point-to-point encryption, and the issues of key distribution, exchange, and the mechanisms used to automate encrypted communication. Everyone seems to know that IPSec will save the world and is the answer to all our security problems (and I have my tongue firmly planted in my cheek), but few know what IPSec really does, from a "features and benefits" point of view. Of particular use and interest are the sections labeled "Deployment Example." These are small case studies that show the technology in action and discuss some of the decisions and processes that came before deployment.

The section covering public key cryptography along with file and message encryption is perhaps shorter than it should be, although much of the groundwork is done earlier in the book. Missing is a "how to" on setting up a public key infrastructure (PKI) for a corporation to use. There are "Product Examples" in this section, but not "Deployment Examples." Perhaps those will have to wait for a second edition, for although this is a lack in the book, there are not many real-life examples from which to choose. Although discussed in theory for years, this is still "leading edge" in the real world. The chapter on Web servers should prove informative and useful to any organization thinking of deploying (or having already deployed) a Web server.

In the chapter entitled "Secure Electronic Mail," the fact that Smith covers *Privacy Enhanced Mail* (PEM) as a technology more than he covers S/MIME is puzzling, but the basics of PEM are useful for discussion, even if PEM as a technology seems to be dead.

Cryptography Is Necessary

The advertisement on the back of the book (not written by the author, of course) states "Here, in one comprehensive, soup-to-nuts book, is the soution for Internet security: modern-day cryptography." Obviously the claim that cryptography is *the* solution for Internet security is way overinflated; modern-day cryptography is not *the* solution, but, cryptography is an important part of a "balanced" security solution. Smith does an admirable job of making this heretofore…well, cryptic… subject, understandable, interesting, and even enjoyable.

—*Frederick M. Avolio, Avolio Consulting,* **fred@avolio.com**

Call for Papers

The Internet Protocol Journal (IPJ) is published quarterly by Cisco Systems. The journal is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. The journal carries tutorial articles ("What is…?"), as well as implementation/operation articles ("How to…"). It provides readers with technology and standardization updates for all levels of the protocol stack and serves as a forum for discussion of all aspects of internetworking.

Topics include, but are not limited to:

- Access and infrastructure technologies such as: ISDN, Gigabit Ethernet, SONET, ATM, xDSL, cable fiber optics, satellite, wireless, and dial systems
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, trouble-shooting, and mapping
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, network computing, and quality of service
- Application and end-user issues such as: e-mail, Web authoring, server technologies and systems, electronic commerce, and application management
- Legal, policy, and regulatory topics such as: copyright, content control, content liability, settlement charges, "modem tax," and trademark disputes in the context of internetworking

In addition to feature-length articles, IPJ contains standardization updates, overviews of leading and bleeding-edge technologies, book reviews, announcements, opinion columns, and letters to the Editor.

Cisco will pay a stipend of US$1000 for published, feature-length articles. Author guidelines are available from Ole Jacobsen, the Editor and Publisher of IPJ, reachable via e-mail at `ole@cisco.com`

Fragments

ICANN

The *Internet Corporation for Assigned Names and Numbers* (ICANN) was incorporated in late October. ICANN is a private, non-profit corporation, managed by an international board, formed to coordinate and administer policies and technical protocols relating to the domain name and address system that permits Internet communications to be routed to the correct person or entity. Its proposed duties include those now performed under U.S. Government contract by the *Internet Assigned Numbers Authority* (IANA), whose Director, Internet pioneer Jon Postel, died on October 16th. ICANN has elected its Initial Board and chosen Michael M. Roberts as its Interim President and Chief Executive Officer. In addition, the Board chose Esther Dyson as its Interim Chairman, and appointed an Executive Committee consisting of Dyson, Gregory L. Crew, Hans Kraaijenbrink and Roberts. The other Initial Board members include Geraldine Capdeboscq (France), George H. Conrades (United States), Gregory L. Crew (Australia), Frank Fitzsimmons (United States), Hans Kraaijenbrink (The Netherlands), Jun Murai (Japan), Eugenio Triana (Spain), and Linda S. Wilson (United States). ICANN was originally proposed by Postel on behalf of a broad coalition of Internet stakeholders in response to the request by the U. S. Government last June that the Internet community create a global consensus non-profit corporation to which the U.S. could transition the responsibility for overseeing and funding those coordination activities. For more information, see:
`http://www.iana.org/index2.html`

APRICOT '99

The *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT) will be held in Singapore, March 1–5, 1999. APRICOT provides a forum for key Internet builders in the region to learn from their peers and other leaders in the Internet community from around the world. The week-long summit consists of seminars, workshops, tutorials, conference sessions, birds-of-a-feather sessions, and other forums—all with the goal of spreading and sharing the knowledge required to operate the Internet within the Asia Pacific region. For more information, see: **`http://www.apricot.net`**

Send us your comments!

We look forward to hearing your comments and suggestions regarding anything you read in this publication. Send us e-mail at: **`ipj@cisco.com`**

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

The Internet Protocol Journal is published quarterly by the Cisco News Publications Group, Cisco Systems, Inc.
www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com

The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive, M/S SJ-J4
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

Bulk Rate Mail
U.S. Postage
PAID
Cisco Systems, Inc.