# Scalable ATM Encryption

*Lyndon G. Pierson*, Sandia National Laboratories, Albuquerque, NM

**ABSTRACT:** *Customers of Asynchronous Transfer Mode (ATM) services may need a variety of data authenticity and privacy assurances. Cryptographic methods can be used to assure authenticity and privacy, but are hard to scale for implementation at high speed. The incorporation of these methods into computer networks can severely impact functionality, reliability, and performance.*

To study these trade-offs, a prototype encryptor/decryptor is under development. This effort is to demonstrate the viability of implementing certain encryption techniques in high speed networks. The research prototype is being designed to process ATM cells in a SONET OC-3 payload. This paper describes the functionality, reliability, security, and performance design trade-offs intended to be investigated with the prototype.

To improve functionality of high speed computer networks, a wide interoperability of applications, network hardware, and network software is desired. This is facilitated by adherence to standards and extending those standards as appropriate. The signalling required to establish customer selected security assurances for traffic in these networks should conform to the evolving ATM and SONET standards.

High end-to-end reliability (availability) requires even higher reliability of intermediate components and automatic fail-over to redundant components where feasible. Cryptosystems which are improperly keyed or which have lost synchronization can ruin the availability of an otherwise reliable communication system. Therefore, key management and synchronization must be carefully implemented so that the availability of encrypted services is as high as the availability of unencrypted services.

End-to-end throughput performance depends on the communication processing rates and signalling rates of network components and also on the delay and error rate to which the traffic is subjected. Clearly, encryption/decryption processes should subject network traffic to as little additional delay and error rate as possible.

Network security protections should be carefully matched to the threats against which protection is desired. Even after eliminating unnecessary protections, the remaining customer-required network security protections can impose severe performance penalties. These penalties (further discussed below) usually involve increased communication processing for authentication or encryption, increased error rate, increased communication delay, and decreased reliability/availability. Protection measures involving encryption should be carefully engineered so as to impose the least performance, reliability, and functionality penalties, while achieving the required security protection.

## End-to-End Encryption vs. Link Encryption

Link encryptors typically encrypt each and every bit of a synchronous communication line at one end of a leased or private circuit, and decrypt each and every bit at the other end. End-to-end encryption can be thought of as occurring at a higher layer of communication protocol, and involves identifying and passing the control information associated with each data packet, while encrypting the payloads of selected data packets. Since the control information is not encrypted, this allows the processing of such packets at intermediate equipment without decryption.

The encryption and decryption processes are usually simpler (and faster) for link encryption than for end-to-end encryption. Link encryptors have fewer decisions to make, and need not identify the beginning, end, or control information of each data packet.

The encryption and decryption processes are intended to be computationally infeasible unless one holds the cryptographic "key". This usually causes the computations to become intensive even on behalf of the authorized "key" holder.

Encryption is computationally intensive, and end-to-end encryption is not only computationally intensive, but also communication processing intensive. This tends to cause the commercial availability of high speed link encryption and end-to-end encryption equipment to lag the availability of high speed communication equipment and service offerings.

High speed communication equipment and service offerings become available as higher speed switching components (faster transistors) become available. In order to keep up with data

Note: This work was supported by the U.S. Department of Energy under contract DE-AC04-76DP00789.

rates available through high speed communication equipment, encryptors typically require a great many of the more expensive, faster switching components. Since the consumer market for encryption has been relatively small, faster encryptors typically become available only when faster transistors become "affordable".

This research is applying parallel processing techniques to scale the speed of encryption processing without requiring the availability of higher speed components. This approach may still require higher speed components to perform parallel-to-serial and serial-to-parallel conversions for transmission, but will require fewer higher speed components throughout the encryption processing function.

## Security

A "dictionary lookup" crypto-analytic attack involves matching the cyphertext of a message against recurring identical cyphertext messages (thereby indicating a subsequent identical plaintext was transmitted). A "playback" attack involves "re-playing" a given cyphertext in order to spoof the repetition of a message. To prevent "dictionary lookup" and/or "playback" attacks, cryptosystems employ various feedback methods. The feedback causes repeated plaintext messages to encrypt into different cyphertexts. Implementation of these feedback methods complicate the scaling of the encryption/decryption process.

## Functionality

End-to-end encryption can provide greater granularity of protection by encrypting simultaneous communication streams with separate sets of cryptovariables. This requires hardware which can quickly switch "cryptovariable contexts" to encrypt cells from multiple concurrent virtual circuits.

For some security purposes, the closer encryption is implemented to the user's application (in terms of communication protocol layers), the better. Other security purposes may demand encryption and/or authentication at other protocol layers. This may result in super-encryption of higher layer encrypted packets by encryption also performed at lower layers. Such "super-encryption" makes cryptographic synchronization loss detection and recovery more difficult for the lower layer encryption processes.

### *Scalability and Interoperability*

Prior to the availability of "variable bit rate" communication services, encryption/decryption processes at both ends of the network were required to operate at the same rate. Link encryptors, encrypting each and every bit of a synchronous communication line, have to decrypt each and every bit at the same rate as encrypted. End-to-end encryptors likewise matched a common media clock rate. With the advent of "variable bit rate" communication services via ATM switchgear, a requirement now emerges for implementations which can inter-communicate, yet operate at different data rates. For example, ATM cells encrypted for transmission into an ATM network via an OC-3 (0.155 Gb/s) or even OC-48 (2.4 Gb/s) interface may be decrypted at a DS3 (0.045 Gb/s) interface. The lower speed, lower cost encryptor/decryptor will implement internal encryption/decryption processes at a lower data rate, likely with a lesser degree of "parallelism".

This requirement, for encryption/decryption processes which can be scaled for high speed, yet interoperate with unscaled or lesser scaled versions, is difficult to achieve in most cryptosystems which use feedback methods to protect against dictionary lookup and replay attacks.

## Reliability

Experience shows that encrypted communications suffer lower reliability than unencrypted communications. One reason is simply the logistics of securely getting encryptor/decryptor pairs keyed with the same key. Murphy's law applies to encryptors: the probability of mismatched keys is great. To overcome this difficulty, asymmetric (public key) key management systems can be employed, but with a corresponding increase in signalling complexity and overhead.

Once properly keyed, cryptosystems which employ feedback to deter dictionary lookup and playback attacks must be synchronized in order to encrypt and decrypt properly. Once synchronized, these cryptosystems remain synchronized until/unless bits are inserted into or deleted from the cyphertext to be decrypted. When this occurs, more or fewer bits are received for decryption than were encrypted. This causes the decryption process to become "out of step" with the encryption process. These "bit count integrity loss" errors, unlike errors which change ones to zeros or vice versa, cause all subsequent plaintext to be improperly decrypted until re-synchronization occurs. Except for a special class of "self-synchronizing" cryptosystems, the communication reliability is a function of how fast and how accurately this "crypto sync loss" state can be detected and recovery achieved.

Most crypto sync loss detection methods involve identifying expected patterns in the decrypted data. These methods work well for low data rate encrypted communications. Because unsynchronized, improperly decrypted data resembles cyphertext (random data), these "pattern matching" methods do not scale well for high data rate communications. The frequency of occurrence of any given pattern in the random unsynchronized output increases with data rate. These detectors become unreliable at high speed because the patterns expected in properly decrypted data are quickly found also in the unsynchronized, improperly decrypted data.

A Sandia developed method (U. S. Patent # 4977596) which measures the randomness of the decrypted data overcomes this problem and scales well for implementation at high data rates.

Once the synchronization loss has been detected at the decryptor, a resynchronization request must be communicated to the encryptor. The encryptor must then send sufficient information to the decryptor to re-establish cryptographic synchronization.

Cryptographic synchronization recovery for these systems requires a period of time to detect sync loss and a round trip delay to request and receive the synchronization information. During this time, encrypted communication bandwidth is unavailable to the user.

While synchronization recovery is taking place, some method of preventing the delivery of “random” improperly decrypted data to end equipment or processes should be implemented. If not, these processes (expecting plaintext) may interpret patterns in the random data as proper control information. This may cause these end processes to transition to unexpected states, so that data processing cannot continue when communication is restored.

## Performance

To achieve high performance, the end-to-end delay and error rate must be minimized. Some encryption/decryption processes can add significant network delay to that experienced by unencrypted traffic. Some encryption/decryption processes decrypt single bit errors in the cyphertext into multiple bit errors in the decrypted plaintext, “magnifying” the error rate to which the cyphertext is subjected.

As network data transfer rates increase, communication becomes less and less efficient due to the increased delay-bandwidth product associated with the end-to-end transactions. This delay-bandwidth product can be thought of as the number of bits in transit which have left the transmitter and have not yet been received by the distant receiver. As this amount of “data in transit” increases, so do difficulties in implementing efficient error control and flow control. The performance of high speed communications systems will be increasingly sensitive to delay and error characteristics.

One of the reasons Asynchronous Transfer Mode (ATM) switching technology is expected to scale to support high speed networks is the low delay incurred as traffic passes through each switching node. Short cell headers and fixed length cells enable ATM switching implementations which achieve switching delays limited only by the time required to assemble the five byte cell header.

Clearly, implementors of encrypted services should strive to minimize additional network delay due to encryption, and to eliminate error magnification by careful choice of encryption method.

## Summary

For encryption/decryption in the high speed “variable bit rate” ATM networking environment, these characteristics are highly desirable:

1. Encryption of multiple identical plaintext messages into differing cyphertexts (to deter dictionary lookup and playback attacks).
2. Scalability of encryption speed (independently of speed of switching elements)
3. Interoperability of scaled and lesser scaled implementations.
4. Little or no error rate magnification.
5. Minimum traffic delay due to encryption.
6. Scalable key management.
7. Fast context switching of cryptovariable context between cell streams (less than header processing time).
8. Fast detection and recovery from cryptographic sync loss.

A “research prototype” encryptor/decryptor is under development. This prototype (not a product) is intended to demonstrate the viability of achieving these objectives by processing ATM cells in a SONET OC-3 payload.

## Bibliography

1. Simmons, Gustavus J. (ed.), Contemporary Cryptography, IEEE Press, New York, 1991.

2. Schneir, Bruce, Applied Cryptography, John Wiley & Sons, New York, 1994.