

Cover sheet

Rex Buddenberg
Code IS/Bu
Naval Postgraduate School
Monterey, Ca 93943

Budden@nps.navy.mil
831/656-3576

Your abstract for the 2002 CCRTS has been accepted and was placed into the IS/IO Track under RADM Gary Wheatley.

Title: Information Security

Information Security

Rex Buddenberg
Code IS/Bu
Naval Postgraduate School
Monterey, Ca 93943
budden@nps.navy.mil
July 1995, rev February 2002, rev April 2002

Abstract.

Security in information systems is a complex problem. Single solutions to complex problems don't exist and matching the appropriate solution (or more accurately, a set of solutions) to a requirement is necessary.

What's in and what's out?

Before we can directly attack the information security problem, we must define terms and scope. This paper addresses a broad spectrum of information security including this list of definitions.

Definitions:

- * Confidentiality. Unintended recipients can't read our traffic. Confidentiality includes secrecy of the data.

- * Authenticity. Unintended originators can't fake traffic. Nobody forged my messages. Authenticity is a superset of integrity.

- * Integrity. Traffic hasn't been tampered with. What you got is what I really sent you.

- * Non-repudiation. I can't get away with saying something and later denying it.

- * Access control. Unauthorized users can't use network and computing resources. More colloquially, keep the riff-raff out of my corner of the 'net.

- * Assurance of service. The network is available for use when I need it. Resistance to denial of service attacks.

- * Traffic analysis. Ability to derive intelligence from the addresses of messages, even if the contents are confidentiality-protected.

- * Traffic flow analysis. Derivation of intelligence inferences by observing flows to and from commands and individuals.

* Interceptability. Ability of unintended recipients to receive traffic (regardless of whether they can read it).

* Jammability. Vulnerability of a link to interruption by signal interference.

Outside of the scope of this note are:

* Availability and survivability issues. Some commentators (including the Defense Science Board, for example) include these issues under the general title of security. This is understandable if you extrapolate from the electric power perspective -- security of the electrical industry clearly includes. I am treating this issue as a legitimate plowshares-->swords issue, but not a security issue. Disaster recovery is a subset of the survivability issue. Resistance to intra-network denial of service attacks is definitely included in the security domain however.

* QoS control. Ability to allocate networking resources to the highest priority needs is a critical plowshares-->swords issue too. But it's not security.

Existing standards framework work.

ISO 7498-2, the security architecture reference model, applies the possibility of providing security services at different layers:

Service	Layer
Confidentiality	1, 2, 3, 4, 7
Authentication	3, 4, 7
Integrity	3, 4, 7
Access control	3, 4, 7
Non-repudiation	7

Unfortunately, this table is typical ISO -- it tells you at which layers security services are theoretically possible. But it doesn't provide any information regarding where it's a good idea and where it's a clumsy and expensive idea that is counter to interoperability objectives. Remember that the Reference Model, including the security part, is ISO's means of deconflicting various committees ... and in this case, not much more. Additionally, security issues such as traffic analysis, traffic flow analysis and jammability are not considered in the ISO model.

Further, the existence of an ISO standard tells us nothing about the existence of any products. If we intend to beat commercial industry plowshares into swords, this is a critical omission of our inquiry if not of the standard. The programmatic landscape is littered with a few attempts, like CANEWARE and BLACKER, that attempted to solve security issues -- principally confidentiality -- at the middle layers. The solutions were clumsy enough never to see the implementation light of day.

Organize thinking.

While the Security Architectural Reference Model is insufficient to our needs, the ISO Reference Model which defines services at different levels of abstraction is a very useful place to start. We can define and classify security services in much the same way that we classify network services. The model provides the means to properly match requirements and solutions. Furthermore, while the ISO Reference Model has been applied to networks, it is equally useful when applied to end systems (e.g. computers) attached to the network.

Let's see if we can add a bit of common sense. Thumb rule: the higher the layer at which you can gain appropriate security service, the less you have to depend on the network to provide the service. For example, with secure e-mail -- an Application Layer implementation of a security service -- all security functionality is provided in end systems, none is required of the network infrastructure (links, routers, gateways, etc). This means that it is not necessary to own or control the network in order to have secure service. In concise terms, with application layer security we have confidentiality and authenticity over untrusted networks. The cost of a security service is going to be proportional to the size of the security enclave that you must secure. Limiting the enclave to a single end system has considerable attractiveness -- indeed when that end system is a single user (e.g. PC), then a great deal of the Orange Book requirements become irrelevant. Which means that the first possibilities to examine are those at Layer 7.

Organizing matrix.

ISO RM hint	Buzz	Problem	Solution	Examples
7, Application	Secure the data	Confidentiality ,Authenticity	Object Level Security	S/Mime, ssh, ssl, VPN
3-4, Network and Transport	Secure the network/ computer (not the data)	Perimeter protection of enclave Prevent DOS attacks	firewalls Intrusion detection, MAC, DAC, TEMPEST, physical	passwords locked doors biometrics
1-2, Physical and Data Link	Protect the pipe	Traffic analysis Tfc flow analysis Jammability Detectability	Link crypto LPI/LPD spread spectrum	KG-84 STUIII wireless LAN

Application Layer security.

The focus at the application layer is on security of the data. The network that it flows through and the computers that it is stored in are irrelevant -- the data itself is the security target.

For example, an e-mail message that has body parts that have been encrypted and digitally signed now has confidentiality and authenticity protection regardless of the means of delivery, whether over a private network, the public Internet, the enemy's network or even sneakernet.

Advantages of this media independence include:

- * since the data committed to the network is encrypted and 'safe', we have no inhibitions against mixing multiple levels of classification on the same communications system. The reason for security-segregated networks disappears.

- * cross-program and cross-ally issues become much more tractable. Sharing network connectivity and sharing data are two separable problems.

- * protecting the data itself removes us from the 'chain as strong as weakest link' situations that inhibit scale and interoperability

Enclave Protection

Enclave protection, in abstract form, is setting up a perimeter and then protecting it from penetration. In networks, we set up these communities of interest and typically protect them with firewalls. In individual computers, we do this by setting up file systems and then allowing certain individuals to access just parts of the file system (e.g. home directories). In databases, certain users may be accorded read or write privileges to certain columns or rows, but not allowed into others. Note that the protections all apply to the perimeter and are focussed on securing the network or computer or database -- not the data. The shortcomings of this means-ends mismatch are well known:

- * These protections are easily circumvented through insider attacks
- * The protections tend to be brittle; a break in any location can compromise a great deal.
- * The protections here are best targeted at avoiding denial of service attacks, not in protecting your data.
- * Mixing multiple levels of security with these mechanisms is very difficult to do -- the grail has been elusive for many years.

Link protection.

Protecting individual links has been the norm in the US military for many years now. It has taken the form of link encryption using devices such as the KG-84 (and it's predecessor the KW-7) and telephone connections using STU-IIIs (secure telephone units). These are useful protections against traffic analysis and traffic flow analysis vulnerabilities. Low Probability of Detection and Intercept schemes: spread spectrum in radio, hard-to-tap fiber optic in wirelines, TEMPEST in electronic equipment, target similar vulnerabilities as well as provide jam resistance. Since jamming is usually only a problem in radio networks (as opposed to fiber optic and copper links), you don't find jam protections in the wired portions of the network. Locks on the computer's power switch and password protection of the BIOS are in this category. Indeed, in general, these protections are good for a single link or single computer only -- they have to be peeled off before the link can feed its contents into a router. This means that link protection is a very poor way to

provide confidentiality since all the routers must be both physically and logically protected and the network can only run at a single level of classification.

Some examples: the good, the bad and the ugly.

Credit card transactions over the Internet.

Most of us have now bought something over the net. When you're ready to buy, your browser opens a secure connection to a server (signified by the closed padlock icon on most browsers), you type in your credit card number and the deal is closed. The requirement here is for confidentiality - - you don't want somebody to crib your credit card number. The solution under the icon is a Secure Sockets Layer connection where the browser and server exchange keys and encrypt the data. So far, so good -- we've matched a confidentiality requirement with an object level security solution (across the top line of the matrix above). And none of the reported theft of credit card numbers appear to involve any breakage here.

But several thefts of credit card numbers have been reported. What happened? Once your credit card has been securely transmitted it is stored by the server ... in plaintext form. Typically, the server is not using object level security to protect the data, but the computer is access-controlled with a password and the read/write privileges associated with access to certain files. In other words, we're seeing perimeter control solutions applied to confidentiality problems. Once those perimeter controls are breached, not just one credit card, but the whole file of them is compromised.

A better approach would be to keep the same object level confidentiality protection on the credit card data itself. In other words, store the encrypted data and only decrypt it when needed. The existing access control mechanism shouldn't be abandoned, but it becomes another layer in the defense.

The Walker Case

In 1985, US counterespionage uncovered an insider attack against the Navy's worldwide communications system. The critical breach involved sale of the KW-7 (crypto machine) key cards to the Soviet Union. This compromise had been going on for years. The security requirements for the Navy's system included confidentiality and authenticity, but also most of the other items in the matrix above, including resistance to traffic analysis. The solution was to use link layer encryption (it was known as bulk encryption then) to solve both. In the form of our matrix, we were using a link level solution to solve an application level problem. This resulted in a brittle system -- once it is compromised in one place (the key cards from a communications station in Philippines or Stockton) it broke everywhere. This is because most KW-7s were keyed alike for interoperability reasons.

The Navy's intelligence community was less heavily impacted by the Walker revelations than the general service Navy communications because the intelligence community used a superencryption scheme known then as Streamliner -- the data was end-to-end encrypted before it was passed to the communications station for transmission. This meant that for this traffic, the traffic analysis

vulnerability became exploitable by the compromise but the confidentiality remained intact. In today's Internet lingo, the Streamliner system is a form of virtual private network.

Ironically, the sea services used what we would today call object level encryption (then known as off-line encryption) prior to the advent of the KW-7. The problem was that the offline systems (Adonis, Diana, etc) were labor intensive. Rather than automate the offline systems (something that computers can easily do today), we shifted targets to securing the communications pipe rather than securing the data and a brittle system resulted.

We must point out that we seem not to have learned the lessons from the Walker case: we replaced the compromised KW-7 with the newer KG-84 but the basic bulk encryption approach persists. We see it today in SIPRNET (Secured IP Router NETwork). The security segregation requires us to plumb the network with both SIPRNET and NIPRNET (with negative availability and performance effects). And because we rely on link level encryption for confidentiality, we have to segregate SIPRNET so far as to deny access to our allies. Media-independent approaches to confidentiality would greatly improve the interoperability cost and speed of installation factors.

Halting steps forward?

At this writing (March 2002) the Department of Defense is circulating a draft "Overarching Wireless Policy."

The draft policy's scope explicitly includes any radio networks to be connected to the Global Information Grid - a fairly broad scope that clearly includes such things as Navy Fleet Satellite Communications and DoD MilStar and their follow-on programs. Since the most appropriate solutions for authenticity and confidentiality are application layer solutions targeted at the data rather than the media, the scope is inevitably larger than just wireless - it must encompass virtually all DoD networks¹.

A policy section sentence states: "Only assured channels employing NSA-approved, Type-1 end-to-end encryption shall be used to transmit CLASSIFIED information". The critical term here is 'end to end'. This policy statement seems to indicate use of application level security for confidentiality purposes as advocated by this paper. This would be a very significant improvement in the military security posture if the interpretation is correct.

Impact. It also represents a significant shift in programs's budgets, putting the emphasis on confidentiality in applications rather than in networks.

One step back. But the definition of 'end to end' is stated in the Definitions enclosure: "End-to-End. AIS from the end user device up to the security border of a DoD network or between two user devices connected by a DoD / non-DoD network (to include the air interface)." This definition is confusing enough to make the object quite unclear. Further, it reads much like a definition of link level security rather than application level.

General observations and conclusion.

¹ Reading the draft, I'm not at all sure this was intentional, but that's the effect.

From these examples, we can draw some general conclusions:

- * The higher on the matrix you can solve a security problem, the better. In particular, if you can solve confidentiality problems at the application layer, you can use general purpose network. This has obvious cost benefits (no special purpose plumbing), and it increases the interoperability quotient -- you can share what networks that do exist with others.

- * None of the solutions are mutually exclusive. It's entirely possible to solve the confidentiality problem with, say, end-to-end secure e-mail, communicate entirely within a closed enclave (carefully firewalled or air-gapped to keep out outsiders) and use link encryption to frustrate traffic analysis by eavesdroppers. Both of the latter two measures help the confidentiality picture -- they're layers in the layered defense -- but they're there primarily for other purposes.

- * When we consider acquiring information systems, we want to express the lower layer requirements to the 'plumbers' -- those who build and provision the network. We want to express the top layer requirements to the application designers. Mixing these signals (graphically visualized as crossing the matrix diagonally) results in asking the right requirements ... but of the wrong providers. Which is almost as bad as specifying the wrong requirements.

- * Most importantly, the specific security requirements must be properly matched with a solution that directly targets the requirement. In the matrix above, this is visually illustrated horizontal lines between problem and solution; diagonal traces indicate a mismatch.