

# A<sup>1</sup> Security Architecture for 802 Networks?

Mick Seaman

This note is a sketch of a potential security architecture for 802 Networks. Its principal purpose is to document the 'obvious'<sup>2</sup>. This first draft is far from complete and is being circulated purely to further the weekly telephone discussions of the Security ECSG. It is written for those who are not security experts by someone else who is not, taking the view that this is a reasonable thing to attempt, and that a security project will not have succeeded unless its results can be explained and implemented<sup>3</sup> by those with little detailed understanding of security.

The proposed architecture is based on 802.1X taken together with the developments proposed to make it more useful and flexible in networks comprising both 802.11 and 802.3 LANs.

---

<sup>1</sup> "A" but not "the", there being some many different threats and perspectives on those threats.

<sup>2</sup> That is it contains little in the way of new ideas, not that it represents the only approach to the problem, a common view of what problem we are trying to solve, or does not need correcting.

<sup>3</sup> Though not designed.

## Why link layer security

The first question that has to be answered is why have link layer security at all? Surely communication is best protected end to end<sup>4</sup>, at the session or application level using IPsec<sup>5,6</sup> or related technology?

Our answers are that:

- 1) The session level communication traverses an intervening network, consuming resources on that network that a network owner might only wish consumed by authorized parties<sup>7</sup>.
- 2) Allowing communicating parties to transmit frames across an intervening network can expose the security weaknesses of the other systems<sup>9</sup> attached to that network to attack. Most common systems have countless points of weakness, some completely accidental, some intentionally introduced to support 'efficiency' or 'plug and play' operation.
- 3) Many LAN protocols that configure the network or the services delivered by the network are not intrinsically secure, and not all these protocols can be redesigned easily to use secure communications. Even if they were so redesigned it is doubtful whether satisfactory implementations would be available soon, or networks already operating could be migrated.

## The nature of a solution

The identified needs (above) suggest the following framework, assumed problem, and security mechanisms. The rest of this note discusses them using 802.1X terminology as convenient.

A network owner creates and manages a secured network, that is a network that has a secure perimeter. Within the secured network no constraints are placed on the communications between stations<sup>10</sup>, but any station external to the secured network (a *Supplicant*) has to petition the owner to be allowed access. Access to the secured network is provided through *Network access ports*, each equipped with *Authenticator* functionality and capable of denying or restricting access. Each Authenticator has little local configuration information but uses the services of an *Authentication Server*, provided and managed by the network owner.

Of course this model, illustrated in Figure 1, was deliberately chosen to be very similar to that used to control dial up internet access<sup>11</sup>, and has the general flavor of a firewall solution<sup>12</sup>.




Figure 1 - Framework

<sup>4</sup> Either bi-directionally, or within a closed group of systems.

<sup>5</sup> When 802.1 began work on .1X there was quite a body of opinion in support of the notion that the only true security was application level security, and that the proposed work was irrelevant given the imminent deployment of comprehensive application level security frameworks.

<sup>6</sup> While this objection is not as common as it was, it does have some merit and is not completely dealt with by the refutation given. In certain scenarios it is possible that protecting only that traffic which can be covered by end to end IPsec provides an appropriate level of security.

<sup>7</sup> Or at least parties that can be charged

<sup>8</sup> These resources may be consumed even if one of the parties supposedly involved in a session wants no part of the communication

<sup>9</sup> These may or may not be readily identifiable as the intended destination of frames transmitted by malicious parties.

<sup>10</sup> At least not by the security mechanisms discussed in this note.

<sup>11</sup> It differs in one very important respect. The supplicant does not have to transmit ordinary data frames in a special encapsulation (such as ppoE) to mimic dial up. Once it is granted unqualified access the network perimeter access is granted

<sup>12</sup> Though good firewalls are more elaborate system constructions.

## A simple beginning

802.1X as first proposed<sup>13</sup> attempted to provide a practical solution to a simple scenario<sup>14</sup>. Ports allowing access to a LAN were to be placed in a semi-public space, e.g. a conference room used by visitors to a company, and access to the LAN would be controlled by a dedicated bridge port taking instructions from an authentication server for the company<sup>15,16</sup>.

From the point of view of getting a result, i.e. producing a working specification, in a finite time this scenario is effective as it admits very few threats:

- 1) The human (let's call him or her 'H') who is attempting to connect to the LAN with a laptop (the supplicant) is presumed (if authenticated and authorized) to be well intentioned, moderately competent, and careful of the LAN and attached resources.
- 2) H can inspect the laptop, provide the physical cable from the laptop to the wall jack, and verify by inspection that no intruder is attached (there is no shared hub half way down the cable).
- 3) The wall jack is firmly attached to the wall, and intruders are not believed to have sufficient uninterrupted access to complete building works to insert equipment behind the jack.
- 4) The physical wiring behind the wall jack is believed to be wired into the bridge port controlling access with semi-permanent point to point wiring using the cable routing commonly used in office walls and cube layouts, so is not readily susceptible to accidental network additions by the company's own non-IT staff.
- 5) The company's own staff, who have direct physical access to the secured LAN, are

presumed not to be directly interested in helping intruders access resources. To put it another way – if they are of mischievous intent the company has much more to worry about than controlling open access wall jacks in meeting rooms, since they can physically access all sorts of equipment in the network directly.

---

<sup>13</sup> It became more complicated or sophisticated (depending on your point of view) fairly rapidly as Radius and EAP were leveraged, but I claim authoritative knowledge of the first proposal to .1, initiated by Vipin Jain. See US patents 6,021,495; 6,311,218; and 6,367,018, for some background and focus on concerns at the time.

<sup>14</sup> It is worth discussing the scenario because one of the obstacles to practical security, at a cost a customer is willing to bear, is excessive generality in the solution. Curiously at the time this generality was advocated as supporting ports on shared media hubs, to lower physical equipment costs. The other major obstacle to getting something done was the view that application level security was the final and only true answer.

<sup>15</sup> Use of LDAP and individual or person object data (I forget the jargon) was high on our list of potential "authentication" methods at the time.

<sup>16</sup> From the outset (and for some time previously) allowing different users with different authorizations to access different VLANs was considered. Since a large number of distinctions could be made and significant apparatus invented for controlling VLAN access invented it was deemed practical to leave this question alone. Given the existence of 802.1X is not now (I hope) at issue, it might be practical to revisit this issue.

## Some additional threats

Amongst the many threats that do not occur in our first simple scenario are:

- 1) Intrusion between the Supplicant and the Authenticator/Network Access Port for:
  - Traffic analysis.
  - Eavesdropping.
  - Replay of data (frames).
  - Modification of (data frames).
  - Insertion of spurious data (frames) on established connections/sessions.
  - Removal of selected data (frames) on established connections/session.
  - Masquerade to establish additional communications/connections/sessions.
  - Theft of credentials for future masquerade.
- 2) Masquerade as a legitimate or expected Authenticator/Network Access Port to deceive the Supplicant into:
  - Revealing confidential data.
  - Revealing credentials that can be used for future masquerade as the supplicant.
  - Accepting data that may be used to attack or compromise the supplicant's system.
  - Communicating to the expected systems through the masquerading Authenticator and thus encountering threats listed under (1) above.
- 3) Masquerade as a legitimate or expected Authenticator/Network Access Port to deceive the Authentication Server<sup>17</sup> into:
  - Facilitating an attack on the supplicant, as per (2) above.
- 4) Intrusion between the Authenticator/Network Access Port and the Authentication Server to:
  - Modify the communication to the Authenticator to support access to the secured network by an unauthorized supplicant.

A further set of threats is omitted from our scenario analysis because the supplicant is not concerned about them, or is taking separate defensive measures:

- 5) Protection of the supplicant against data from the expected<sup>18</sup> secured network that attempts to attack or compromise the supplicant's system.

---

<sup>17</sup> A direct threat to data within the secured network is only possible if the intruding/masquerading system is attached to the network, so many threats that might be expected to occur under this heading do not appear. Once surreptitiously attached to the secured network an intruder can do many evil things.

<sup>18</sup> "expected" in the sense that it is the network that the supplicant intends to connect to.

## Simple scenario benefits

The simple scenario described has a simple solution with some nice properties that we would do well to retain in more complex scenarios when and if possible:

- 1) Communications on the secured network do not have to be encrypted or otherwise protected. Most systems function exactly as they would without security. Existing frame formats are not modified.
- 2) Communication between the Supplicant and the Authenticator does not have to be protected, and frame formats are not modified.
- 3) Communication between the Authenticator and the Authenticating System does not have to be protected, and frame formats are not modified<sup>19</sup>.
- 4) The behavior of the Authenticator can be made very simple as it can act on simple "Authorized"/"Unauthorized" commands sent in clear from the Authenticating System and otherwise simply has to distinguish the authentication/authorization dialogue that it relays in Unauthorized state from the other frames that it must discard<sup>20</sup>.

---

<sup>19</sup> The Authenticating System is within the perimeter of the secured network, as shown in Figure-1. A modest change to this scenario uses a secure tunnel beginning within the secured network to securely convey communication to and from an Authenticating System outside the perimeter of the secured network.

<sup>20</sup> Sadly the initial 802.1X Authenticator became rather more complex than this ideal, which has it transparently passing frames, just snooping on authorization results.

## Mutual authentication

It is possible, even likely, that the system with supplicant functionality will wish to authenticate and authorize the network access port it is trying to use<sup>21</sup>.

Since the security architecture is designed to protect against unauthorized access to the resources that a network access port is safeguarding, each Authenticator may choose its Authentication Server independently, as in Figure 2. To put this another way : just because your boss says I can use his network, doesn't mean that my boss will let you have access to hers.

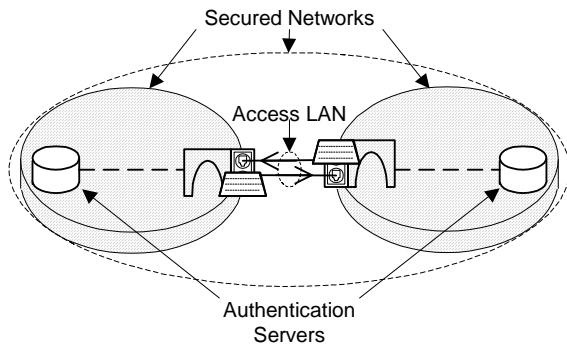


Figure 2

## A growing network

If an authorized network access port is allowed to relay all data traffic<sup>22</sup>, it effectively includes the supplicant within the perimeter of the secured network. If the supplicant is associated with a Bridge Port, with other ports on that bridge providing network access ports then the perimeter of the secured network can grow, as Figure 3 illustrates.




Figure 3

In this scenario it is likely that both bridges, the one initially at the edge of the secured network and the one seeking to participate in that network, will wish to authenticate each other so that neither is accepting frames from or sending frames to an unauthorized party. The bridge that is closest to the authentication server controls the path to that server, and has to authenticate and authorize the other before the latter can proceed. Fortunately no special sequencing is required in the protocol, just suitably long retry timers and counters.

<sup>21</sup> In the simple scenario this was unnecessary since H knows, or at any rate thinks he knows, which network he is trying to access, and trusts that is not malicious.

<sup>22</sup> Or almost all, given that there is a role for some filtering in a network that has no aspirations to security.

## Insecure access link scenarios

The three party framework proposed (Supplicant / Authenticator / Authenticating System) and the simple scenario described above can be used as the starting point for a whole range of scenarios depending on the additional threats each attempts to address.

Scenarios that just include threats from the first two categories, i.e. intrusion between Supplicant and Authenticator, or masquerade as an Authenticator (to the Supplicant), form an interesting subset, illustrated by Figure 4.




Figure 4

In pure form these scenarios imply that data and control frames need not be encrypted or protected within the secured network. The protective security measures introduced to secure the access link can be custom designed to fit the threats admitted for a particular access technology.

## Concatenated secured links

The

*This is it so far, sections to be written:*

- Describe a secured network comprising a set of links each individually secured with mechanisms sufficient for threats to that technology
- Contrast that with an approach where the union of threats (including intruding bridges/network access points) is considered and protection is applied over a number/region of the links in one operation
- Describe a scenario (a sea of troubles, or SoT) where Authenticator to Authenticating System communication is admitted open to compromise/intrusion, either because this is so vital the secured network is not to be trusted, or because the routing of the communication may be in doubt even if the secured network is completely secure, or because good conservative design suggests such an approach
- Consider the relative difficulty of providing supplicant authorization to the authenticator under the "sea of troubles" assumption as opposed to the fully secured backend communication
- Describe the scenario where the whole network is wireless (an SoT variant, but more challenging since it extends

beyond securing A to AS communication)

- Look at the types of system playing Supplicant and Authenticator roles to formulate a sense of the balance between (a) specifying/implementing a single SoT protocol solution and (b) picking two or more threat scenarios to make some cases easier (one of the cases being .1X as of today)
- Augment the model by introducing separate planes in the secured network that are used to segregate traffic. Given that systems within the secured network can be trusted not to couple the planes, the planes themselves do not have to be securely implemented, it's just the behavior of the network access points that does. VLANs (with qualifications) being one way of implementing such a traffic segregation facility.
- Consider the parameters that a network access port needs to be able to place all or part of the incoming traffic on (an) appropriate plane(s), including constraints on how much traffic ought to be admitted (at what priority). Discusses the use of Radius (or Diameter) for distributing these parameters in a way that is tied to authorization, as well as some alternatives.
- Further augment the model to include supplicants roaming between network access points.
- Provide a brief (and not necessarily accurate) introduction to Needham-Schroder methods just to summarize the idea, and support a model of roaming.
- Consider what ought to be done to be honest with ourselves as we bend and flex 802.1X to behave in ways that are not apparent from the current documentation.

Enough for now.