Advertising authentication requirements

Mick Seaman

This note follows up 802.1 Security TG January 2008 discussions. The intent is to work through issues, facilitating both comment prior to the task group ballot of the next draft and discussion at the next meeting.

This note describes the use of controls and advertised requirements and in various scenarios. A very large number of scenarios is possible, corresponding to the combined effects of different settings for the interacting participants. The way in which the scenarios are described is intended to facilitate easy addition of others of particular interest.

While the case for advertising the authentication requirements of an access point seems clear, that for advertising the capabilities of an accessing host (or providing them in an EAPOL-Start) seems less so. The main need seems to come from the desire to provide higher layer authentication as a ‘fallback’ only after an EAP attempt. While a host advertising what it could do seemed a very good idea, it may be that the amount of information really required for ‘capabilities’ is very modest. See particularly Section 6 of this note.

1. Introduction

Access control can delay access to a network: a particularly annoying outcome if one or other of the communicating parties does not support the access control at all or lacks the capabilities that the other would require to grant anything other than the most limited access, subject to the same filters or traffic deflection that would be imposed if no access control protocol were to be attempted.

An Authenticator can reduce these delays by advertising (in an EAPOL-Advertisement) its requirements, and its strategy(ies) for granting access if the Supplicant fails to meet those requirements. A Supplicant might hasten the use of such a strategy by communicating its capabilities in an EAPOL-Start or an EAPOL-Advertisement.

The point of having a strategy for granting restricted or fallback access, rather than granting access immediately, is to avoid having the Controlled Port operUp unnecessarily: hence avoiding the churn associated with acquiring an IP address and operating other protocols only to redo them once secured authenticated access has been provided. Either (or any) of the systems seeking to gain or provide access to or from a LAN can defer access. Useful connectivity will only result when both decide to enable their Controlled Port[1].

There are benefits to knowing both the peer system's capabilities and strategy. If the strategy included waiting for a protocol to produce a result, and that protocol is not implemented, then the associated timeout can be skipped. If one participant attempts the use of a higher layer protocol while the other is still discarding frames received at its disabled Controlled Port, that protocol may not recover promptly when the latter is enabled[2].

Neither EAPOL-Advertisements nor EAPOL-Starts are secure, so a recipient can ignore any capability or strategy data they contain. Any action taken on the data should present no more than a minor irritation—a delay in establishing communication for example—if the frame was actually sent by an attacker instead of a peer that could be authenticated and secured[3]. The data is thus a 'hint' to the peer, and if it is to be relied upon after communication has been secured it needs to be resent securely—either using the MKA Transport with a defined parameter set or using a protocol that is supported by Controlled Port communication.

2. Summary

This note includes:

- A recap and update[4] of the authentication requirements information that can be included in advertisements, and of the Logon Process controls that contribute to the content of those advertisements and guide the behavior of the recipient(s). (3)
- Scenarios between two 'quiet' participants, i.e. neither advertises information, and the outcome of authentication exchanges is guided by each participants undisclosed logon process controls. (4)
- Scenarios when an access point advertises information and an accessing host takes note of it. (5)
- Considerations that suggest that it might be useful for the accessing host to indicate its capabilities to an access point prior to authentication, in order to avoid unnecessary time-outs. (6)

A full understanding of the intent behind the access point's and host's configuration requires knowledge of the authorization controls (PVIDs, filters and other restrictions) applied at any time. These are outside the scope of P802.1af, but their probable intent (as independently applied in each of the communicating systems) is noted for each scenario.

The word 'requirement' is generally used in this note to describe what an access point is capable of or desires in authentication[5], while 'capability' is used of a peer system connecting to that access point. This distinction fits well with EAP Authenticator (requirement) and Supplicant (capability) roles, and provides, but is recognized as a little artificial. This note does not depend on assigning any particular meaning to what are simply a pair of convenient labels.

---

[1] With the following exception: P802.1af D1.8 clause 7.5.3 describes the use of a separate port with selective relay (e.g. of WoL frames) system level mechanisms for providing connectivity to unauthenticated systems for a multi-access LAN.

[2] Consider, for example, a DHCP Client attempting DHCP to a DHCP Server located behind (rather within) an access point. The client's initial attempts to acquire an IP address could be discarded, but the server will not receive any special indication when the access point's Controlled Port is enabled.

[3] Of courses repetitive sending of EAPOL Starts can prevent an EAP authentication from succeeding.

[4] Based on P802.1af D1.8 and taking into account discussion at the January 2008 interim and (possibly) observations made while preparing this note, both of which I intend to put into the next draft of P802.1af.

[5] And desires of subsequent data transfer protection.

3. Authentication requirements

Table 1-1 is an updated version[1] of P802.1af D1.8.

Table 1-1—Authentication requirements information

Octet	Bit	Indicates
1 Immediate full access combinations	1 (l.s.b)	Unauthenticated access
	2	EAP
	3	EAP + MKA
	4	EAP + MKA + MACsec
	5	MKA
	6	MKA + MACsec
	7	Higher layer authentication
	8 (m.s.b)	Vendor specific
2:bits1-3 Restricted unauthenticated access	1 (l.s.b)	Restricted unauthenticated access provided
	2	EAP attempt required for unauthenticated access
	3	MKA attempt/timeout required for unauthenticated access
2: bit 4	4	Restricted, authenticated, unsecured access permitted
2: bits 5-7 Fallback to higher layer authentication	5	Fallback to higher layer authentication provided
	6	EAP attempted required for fallback
	7	MKA attempt/timeout required for fallback.
2:bit 8	8 (m.s.b)	Reserved for future standardization.

These requirements are advertised for each NID, (though PbNAC mechanisms do not distinguish data frames by NID when access is provided[2]) by a port that is part of an access point, i.e. that bridges or route frames. Advertised requirements reflect the administrators knowledge of the following: the whole network; configuration of AAA servers; Radius attributes that those servers provide to the access point following success of an authentication or termination of an authenticated session; the capabilities of the access point; the values of its Logon Process controls; and whether its EAP Authenticator and MKA Entity are enabled. The last two of these comprise:

— logon: Set if the Logon Process is to use results obtained, or to be obtained, from PACP ....

— useEAP: ... when to behave as an EAP Supplicant , if ... logon is set, ...:

  Immediate: ... concurrently with ... MKA with any cached CAK(s).

  MKAfail: Not until MKA has failed, if a prior CAK has been cached.

— unauthAllowed: ... when ... to provide unauthenticated connectivity ... :

  Never: Never.

  Immediate: Immediately, ...

  LoggedOff: Only if logon is not set.

  EAPfail: Not until ... attempt ... using EAP, ...

  MKAfail: Not until MKA attempted.. .

  EAPMKAfail: After attempts ... both EAP and MKA ... .

— unsecureAllowed: ... when to provide authenticated but unsecured connectivity ... :

  Never: Never.

  Immediate: ... when authentication succeeds.

  MKAfail: Not until MKA has failed ... .

  MKAserver: ... if directed by the MKA server.

— auth.enabled: True/False

— supp.enabled: True/False

— mka.enabled: True/False

[1]I believe we agreed to replace the description of Octet 1 bit 1, which was “Open access” with “Unauthenticated access”, and to increase the conditions for fallback to higher layer authentication to match those for restricted unauthenticated access.

[2]Since the NID really represents a network service, and not a VLAN (the latter is a mechanism that may subset the number of accessible network services) it is not possible to say whether the unauthenticated access provide is to the service represented by one NID or another without knowledge of the whole network and all the higher layer protocols.

4. Quiet participants

Logon Process controls provide flexibility when neither requirements nor capabilities are advertised. Assuming a host and an access point both implement the Logon Process controls, Table 1-3 shows some scenarios with the following table entries:

Y condition or option clear, TRUE, or selected

N condition or option clear, FALSE, or not selected)

Table 1-2—Quiet participants

	Scenarios		1	2	3	4	5	6
Access point Controls	supp.enabled		N	N	N			
	auth.enabled		Y	Y	Y			
	mka.enabled		Y	Y	Y			
	logon		Y	Y	Y			
	useEAP	Immediate	–	–	–			
		MKAfail	–	–	–			
	unAuth	Never						
		Immediate	Y	Y				
		LoggedOff						
		EAPfail						
		MKAfail						
		EAPMKAfail			Y			
	unsecured	Never						
		Immediate	Y					
		MKAfail		Y	Y			
		MKAserver						
Host Controls	supp.enabled		Y	Y	Y			
	auth.enabled		N	N	N			
	mka.enabled		N	N	Y			
		logon	Y	Y	Y			
	useEAP	Immediate	Y	Y				
		MKAfail			Y			
	unAuth	Never						
		Immediate						
		LoggedOff						
		EAPfail	Y	Y				
		MKAfail						
		EAPMKAfail			Y			
	unsecured	Never						
		Immediate	Y	Y				
		MKAfail			Y			
		MKAserver						

1: The access point provides both unauthenticated and authenticated connectivity, possibly with a different PVID[1]. However the host will not settle for unauthenticated connectivity unless it has at least attempted EAP (or has no prospect of carrying out EAP because it lacks an appropriate credential, and has no immediate prospect of obtaining one) and currently has a credential that appears suitable. If the EAP attempt succeeds the host will enable its Controlled Port immediately, without attempting to use MKA or MACsec.

2: A similar scenario, with identical host configuration, but the access point would prefer to use MACsec and will only settle for unsecured connectivity when its attempt to use MKA times out, while the host will not attempt MKA but enable its Controlled Port immediately (and thus have no connectivity until the access point times out MKA).

3: The host has MKA enabled, and will attempt to use a cached CAK if one is available, and only use EAP (as a Supplicant) when that attempt fails. If there is no cached CAK it will use EAP immediately, and then use the EAP results with MKA, only settling for unauthenticated connectivity if attempting to use EAP fails. Similarly the access point will not provide unAuthenticated connectivity until an EAP attempt fails (or is timed out). The unAuthenticated connectivity may be set up using a PVID that provides access to a remediation service that depends on some of the parameters used in the EAP attempt.

[1]The ‘PVID’ for a port selects the VLAN used for frames received untagged.

5. Quiet supplicants

Table 1-3 shows some scenarios with a network access point that advertises its requirements, and a host that takes note of these but does not advertise its capabilities. The access point always has supp.enabled

Table 1-3—Quiet supplicants

	Scenarios		1	2	3	4	5	6
Access point advertisement	Immediate full access combinations	Unauth.	Y	–	N	Y	Y	
		EAP	N		Y	Y	N	
		EAP + MKA	N		N	N	N	
		EAP + MKA + MACsec	N		Y	N	N	
		MKA	N		N	N	N	
		MKA + MACsec	N		Y	N	N	
		Higher layer auth.	–	N	N	N	Y	
		Vendor specific	N	N	N	N	N	
	Restrctd. unauth.	Restricted unauth.	Y	Y	Y	Y	N	
		EAP attempt reqd.	N	N	N	Y	N	
		MKA attempt reqd.	N	N	N	N	N	
	Restricted unsecured		N		N	N	N	
	Fallback.	Fallback.	N		N	N	N	
		EAP attemp reqd.	N		N	N	N	
		MKA attempt reqd.	N		N	N	N	
Host Controls	supp.enabled		Y	Y	Y			
	auth.enabled		N	N	N			
	mka.enabled		N	N	Y			
		logon	Y	Y	Y	Y	Y	
	use EAP	Immediate	Y	Y			Y	
		MKAfail			Y	Y		
	unAuth Allowed	Never						
		Immediate						
		LoggedOff				N		
		EAPfail	Y	Y		Y	Y	
		MKAfail						
		EAPMKAfail			Y			
	unsecured Allowed	Never						
		Immediate	Y	Y			Y	
		MKAfail			Y			
		MKAserver						

false: its other control settings should be obvious from the advertisement.

1: The access point does not support any form of authentication, however the host does *not* have unAuthAllowed == Never, so it connects immediately

2: The host has no logon credentials (user/password, whatever), but will enable its Controlled Port so that it can be managed. The access point allows that management. This can be appropriate on a corporate network, so the hosts can be managed out of hours, while there is no doubt that they are in some way permanently associated with the network (or will be treated as if they are).

3: For ‘full access’ the access point requires either authenticated connectivity (using EAP) or authenticated and secured connectivity (using EAP followed by MKA and MACsec, or MKA with a cached key and MACsec). However ‘restricted’ access is available at any time. The host has suitable credentials for the NID (or can attempt to acquire them with what it considers to be an insignificant delay). If the host has a cached CAK for a Key Management Domain advertised by the access point it will attempt to use that. If its MKA times out while attempting that CAK, it will attempt an EAP exchange (as a Supplicant). If that EAP exchange fails it will enable its Controlled Port, with appropriately restricted authorization. However if the EAP exchange succeeds, the host will then use the EAP results for a further MKA attempt, using that MKA to get MACsec operational and then enable its Controlled Port.

4: The access point offers ‘full’ unauthenticated connectivity, authenticated connectivity with EAP, and restricted unauthenticated connectivity after an EAP attempt. These are presumed to be different as follows. The authenticated connectivity is less restrictive than that ordinarily available as ‘unauthenticated’ whereas the restrictive unauthenticated may offer remedial services (possibly those appropriate to the credentials used in the EAP attempt). These changes in the service provided by the access point can correspond to policy controls downloaded as Radius attributes, or to changes in the port’s PVID—again probably as a result of changing a Radius attribute. The host (presumed EAP capable) will disable its Controlled Port until it has at least attempted EAP, and then enable it with authorization appropriate to the EAP result.

5: The access point offers both unauthenticated connectivity and (a path to) higher layer authentication. These are probably both provided on a single VLAN, with a PVID that causes traffic for the port to be bridged on or off that VLAN. The host would have attempted EAP, but seeing that it cannot succeed considers it to have failed, and immediately enables its Controlled Port thus giving its user a path to the higher layer authentication server or servers.

6. Advertising capabilities (or not)

As previously mentionned, there can be a cost to providing unauthenticated connectivity too early in the authentication dialogue. Protocols are run that may have to be rerun once authentication has succeeded, or once failure data is available. Similarly there can be a cost to providing connectivity too late. Initial protocol transmissions by the peer can be discarded on receipt, and when connectivity is provided retries can be infrequent, or even require intervention by the human user of the peer.

There are two possible approaches to dealing with this problem. One is for an accessing host to declare its capabilities and strategies for accepting possibly restricted connectivity: with both systems sharing data it is more likely that the correct connectivity will be provided at the correct time. The other is for the access point to provide restricted connectivity immediately or early, while the accessing host does so later. This latter strategy depends on the fact that the protocols most susceptible to initial packet loss are those initiated by an accessing host, such as DHCP requests. Moreover restarting a host's connection is much more likely to annoy someone.

If we decide not to provide mechanisms to advertise capabilities, an access point that does not know whether an accessing host is listening to or expects advertisements and intends to provide restricted unauthenticated connectivity should not make that connectivity contingent upon an EAP or MKA attempt, i.e. should have its relevant Logon Process controls set as follows:

Access point:

— unauthAllowed: ... when ... to provide unauthenticated connectivity ... :

Immediate: Immediately, ...

— unsecureAllowed: ... when to provide authenticated but unsecured connectivity ... :

Immediate: ... when authentication succeeds.

and let changes (following successful authentication) in PVID or other authorization level mechanisms, such as filters, do the work. The host should have its controls set to match its desires, possibly as follows:

Host:

— unauthAllowed: ... when ... to provide unauthenticated connectivity ... :

LoggedOff:

or

EAPfail:

— unsecureAllowed:

Immediate: ... when authentication suceeds.

The problems with this simple strategy come when both higher layer authentication and remedial restricted unauthenticated connectivity (following failure of EAP or MKA attempts) are to be provided, and support of the latter is to be carried out by placing the accessing host on a different VLAN. In the absence of some capability information from the host the access point is forced to choose between these: the two alternative advertisements in Table 1-4 illustrate this configuration choice ('Fallback' always means 'fallback to higher layer authentication').

Table 1-4—Higher layer or remedial service?

	Scenarios		1	2	3	4
Access point advertisement	Immediate full access combinations	Unauth.	N	N	N	N
		EAP	Y	Y	Y	Y
		EAP + MKA				
		EAP + MKA + MACsec				
		MKA				
		MKA + MACsec				
		Higher layer auth.	N	Y	N	N
		Vendor specific	N	N	N	N
	Restrctd. unauth.	Restricted unauth.	Y	Y	N	Y
		EAP attempt reqd.	N	Y	–	Y
		MKA attempt reqd.	N	N	–	N
	Restricted unsecured					
	Fallback.	Fallback.	N	N	Y	Y
		EAP attempt reqd.	N	N	Y	N
		MKA attempt reqd.	N	N	N	N

1: The access point expects EAP to be used. The accessing host will be given restricted remedial connectivity if it fails EAP and if simply goes ahead and does not attempt to use EAP.

2: The access point expects EAP or higher layer auth.

3: The access point expects EAP, but will fall back to higher layer authentication if EAP fails.

If the accessing host is not EAP capable it could advertise that fact, and in scenario 3 could move directly to using higher layer auth without having to bungle an EAP attempt. But in that direct move is desirable why does the access point not provide scenario 2, possibly advertising that as scenario 4?

Another thing that would be useful, in an EAPOL-Start, would be for the Supplicant to indicate whether it has implemented and enable MKA for use with the EAP results, as that would save an access point flailing

connectivity by first enabling the Controlled Port with unsecured connectivity and then securing the connectivity with MKA.