Extensible Authentication Protocol
Internet-Draft
Expires: July 10, 2004

J. Arkko
Ericsson
B. Aboba, Eds.
Microsoft

January 10, 2004

Network Discovery and Selection Problem
draft-ietf-eap-netsel-problem-00

Status of this Memo

Copyright Notice

Abstract

The so called network discovery and selection problem affects network access, particularly in the presence of multiple available wireless accesses and roaming. This problem has been the subject of discussions in various standards bodies. This document summarizes the discussion held about this problem in the Extensible Authentication Protocol (EAP) working group at the IETF. The problem is defined and divided into subproblems, and some constraints for possible solutions are outlined. The document presents also some existing mechanisms which help solve at least parts of the problem, and gives some suggestions on how to proceed for the rest.

Table of Contents

1. Introduction 3
2. Problem Definition 4
2.1 Access Network Discovery 4
2.1.1 Issues with Access Network Discovery 5
2.2 Identity selection 5
2.3 AAA routing 7
2.3.1 Issues with AAA Routing 8
2.4 Payload Routing 10
2.4.1 Issues with Payload Routing 10
2.5 Discovery, Decision, and Selection 11
3. Design Constrains 13
4. Existing Work 14
4.1 IETF 14
4.2 IEEE 15
4.3 3GPP 16
4.4 Other 17
5. Conclusions 18
6. Security Considerations 21
Normative References 22
Informative References 23
Authors' Addresses 24
A. Contributors 25
Intellectual Property and Copyright Statements 26

1. Introduction

The so called network discovery and selection problem affects network access and wireless access networks in particular. This problem comes relevant when any of the following conditions are true:

- There is more than one available network attachment point, and the different points may have different characteristics.
- The user has multiple sets of credentials. For instance, the user could have one set of credentials from a public service provider and another set from his company.
- There is more than one way to provide roaming between the access and home network, and service parameters or pricing differs between them. For instance, the access network could have both a direct relationship with the home network and an indirect relationship through a roaming consortium.
- The roaming relationships between access and home networks are so complicated that current AAA protocols can not route the requests to the home network unaided, just based on the domain in the given Network Access Identifier (NAI) [4].
- Payload packets get routed or tunneled differently, based on which particular roaming relationship variation is used. This may have an impact on the available services or their pricing.
- Providers share the same infrastructure, such as wireless access points.

The network discovery and selection problem spans multiple protocol layers and has been the subject of discussions in IETF, 3GPP, and IEEE 802.11. This document summarizes the discussion held about this problem in the Extensible Authentication Protocol working group at

IETF.

In Section 2 the problem is defined and divided into subproblems, and some constraints for possible solutions are outlined in Section 3. Section 4 discusses existing mechanisms which help solve at least parts of the problem. Section 5 gives some suggestions on how to proceed for the rest.

2. Problem Definition

There are four somewhat orthogonal problems being discussed under the rubric of "network discovery and selection".

- First, there is the problem of "Access Network discovery". This is the problem of discovering access networks available in the vicinity, and the points of presence (POPs) associated with those networks.

- Second, there is the problem of "Identifier selection". This is the problem of selecting which identity (and credentials) to use to authenticate to a given POP.

- Thirdly, there is the problem of "AAA routing" which involves figuring out how to route the authentication conversation originating from the selected identity back to the home realm.

- Finally, there is the the problem of "Payload routing" which involves figuring how the payload packets are routed, where more advanced mechanisms than destination-based routing is needed.

Alternatively, the problem can be divided to the discovery, decision, and the selection components. The AAA routing problem, for instance, involves all components: discovery (which mediating networks are available?), decision (choose the "best" one), and selection (this is the chosen network) components.

2.1 Access Network Discovery

The Access Network Discovery problem has been extensively studied, see for instance the results of the IETF Seamoby WG, IEEE specifications on 802.11 wireless LAN beaconing and probing process, studies (such as [29]) on the effectiveness of these mechanisms, and so on.

Traditionally, the problem of discovering available networks has been handled as a part of the link layer attachment procedures, or through out-of-band mechanisms.

RFC 2194 [3] describes the pre-provisioning of dialup roaming clients, which typically included a list of potential phone numbers, updated by the provider(s) with which the client had a contractual relationship. RFC 3017 [8] describes the IETF Proposed Standard for the Roaming Access XML DTD. This covers not only the attributes of the Points of Presence (POPs) and Internet Service Providers (ISPs), but also hints on the appropriate NAI to be used with a particular POP. The RFC supports dial-in and X.25 access, but has extensible

address and media type fields.

In IEEE 802.11 WLANs, the Beacon/Probe Request/Response mechanism provides a way for Stations to discover Access Points (APs), as well as the capabilities of those APs. Among the Information Elements (IEs) included within the Beacon and Probe Response is the SSID, a non-unique identifier of the network to which an Access Point is attached. By combining network identification along with capabilities discovery, the Beacon/Probe facility provides the information required for both network discovery and roaming decisions within a single mechanism.

2.1.1 Issues with Access Network Discovery

As noted in [28], the IEEE 802.11 Beacon mechanism does not scale well; with a Beacon interval of 100ms, and 10 APs in the vicinity, approximately 32 percent of an 802.11b AP's capacity is used for beacon transmission. In addition, since Beacon/Probe Response frames are sent by each AP over the wireless medium, stations can only discover APs within range, which implies substantial coverage overlap for roaming to occur without interruption.

A number of enhancements have been proposed to the Beacon/Probe Response mechanism in order to improve scalability and roaming performance. These include allowing APs to announce capabilities of neighbor APs as well as their own, as proposed in IEEE 802.11k; propagation of discovery information over IP, as handled in the CARD protocol developed within the IETF SEAMOBY WG, etc.

Typically scalability enhancement mechanisms attempt to get around Beacon/Probe Response restrictions by sending advertisements at the higher rates which are enabled one the station has associated with an AP. Since these mechanisms run over IP, they can utilize IP facilities such as fragmentation, which the link layer mechanisms may not always be able to do. For instance, in IEEE 802.11, Beacon frames cannot use fragmentation because they are multicast frames, and multicast frames are not acknowledged in 802.11.

2.2 Identity selection

As networks proliferate, it becomes more and more likely that a given EAP peer may have multiple identities and credential sets, available for use in different situations. For example, the EAP peer may have an account with one or more Public WLAN providers; a corporate WLAN; one or more wireless WAN providers. As a result, the EAP peer has to decide which credential set to use when presented with a given set of potential EAP authenticators.

Figure 1 illustrates a situation where the user does not know whether he is connected to access network 1, which only serves the ISP, access network 3, which only serves the company, or access network 2, which serves both.

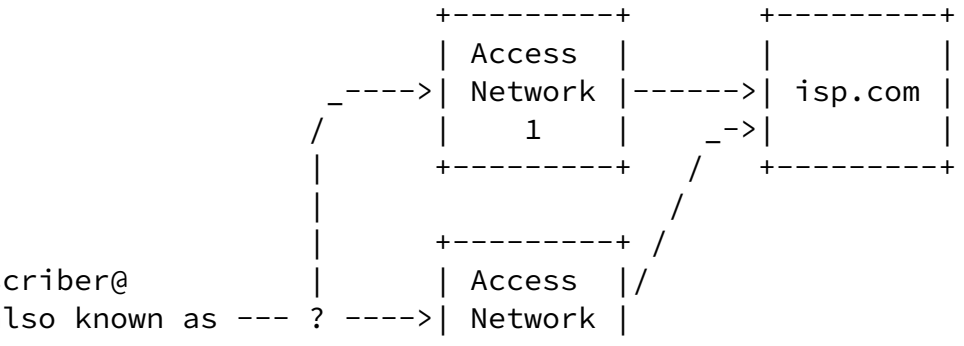

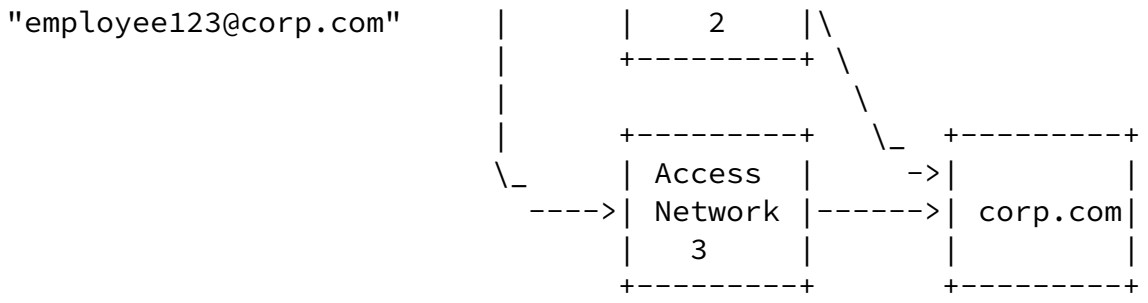

Figure 1: Two credentials, three possible access links

Traditionally, hints useful in identity selection have also been provided out-of-band. For example, via the RFC 3017 XML DTD [8], a client can select between potential POPs, and then based on information provided in the DTD, determine the appropriate NAI to use with the selected POP.

Perhaps the most typical case is a link layer that provides some information about the network before EAP is initiated. For instance, in IEEE 802.11 provides the SSID, though in some cases the client may not learn about all the SSIDs supported by the given access point. In IKEv2 [15], the identity of the responder (typically the security gateway) is provided in the same packet as the EAP Identity Request is transported. In order to use this information in deciding the right identity to use, the provided information has to either match with one of the client's home networks, or the client has to have some other knowledge that enables to link the advertised network and the home network. For instance, the client may be aware that his home network has a roaming contract with a given access network.

It is also possible for hints to be embedded within credentials. In [11], usage hints are provided within certificates used for wireless

authentication. This involves extending the client's certificate to include the SSIDs with which the certificate can be used.

Finally, some EAP implementations use the space after the NUL character in an EAP Identity Request to communicate some parameters relating to the network requesting EAP authentication. However, there is no standard interpretation of the field beyond the NUL

character.

2.3 AAA routing

Once the identity has been selected, it is necessary for the authentication conversation to be routed back to the home realm. This is typically done today through the use of the Network Access Identifier (NAI), RFC 2486 [4], and the ability of the AAA network to route requests to the domain indicated in the NAI.

Within the IETF ROAMOPS WG, a number of additional approaches were considered for this, including source routing techniques based on the NAI, and techniques relying on the AAA infrastructure. Given the relative simplicity of the roaming implementations described in RFC 2194 [3], static routing mechanisms appeared adequate for the task and it was not deemed necessary to develop dynamic routing protocols.

As noted in RFC 2607 [5], RADIUS proxies are deployed not only for routing purposes, but also to mask a number of inadequacies in the RADIUS protocol design, such as the lack of standardized retransmission behavior and the need for shared secret provisioning.

By removing many of the protocol inadequacies, introducing new AAA agent types such as Redirects, providing support for certificate-based authentication as well as inter and intra-domain service discovery, Diameter allows a NAS to directly open a Diameter connection to the home realm without having to utilize a network of proxies. For instance, the Redirect feature could be used to provide a centralized routing function for AAA, without having to know all home network names in all access networks.

This is somewhat analogous to the evolution of email delivery. Prior to the widespread proliferation of the Internet, it was necessary to gateway between SMTP-based mail systems and alternative delivery technologies, such as UUCP and FidoNet, and email-address based source-routing was used to handle this. However, as mail could increasingly be delivered directly, the use of source routing disappeared.

As with the selection of certificates by stations, a Diameter client wishing to authenticate with a Diameter server may have a choice of

available certificates, and therefore it may need to choose between them.

2.3.1 Issues with AAA Routing

No dynamic routing protocols are in use in AAA infrastructure today. This implies that there has to be a device (such as a proxy) within the access network that knows how to route to different domains, even if they are further than one hop away, as shown in Figure 2. In this figure, the user "joe@corp3.com" has to be authenticated through ISP 2, since the domain "corp3.com" is served by it.

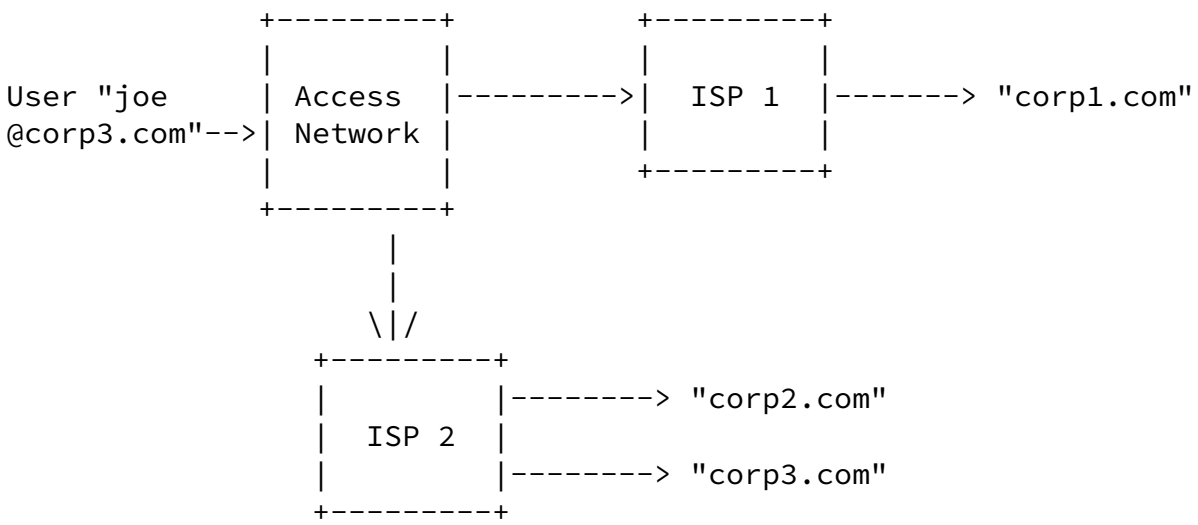

Figure 2: AAA routing problem

A related issue is that the roaming relationship graph may have ambiguous routes, as shown in Figure 3. As billing is based on AAA and pricing may be based on the used intermediaries, it is necessary to select which route is used. For instance, in Figure 3, access through the roaming group 1 may be cheaper, than if roaming group 2 is used. For commercial reasons, intermediaries want to participate the AAA transaction.

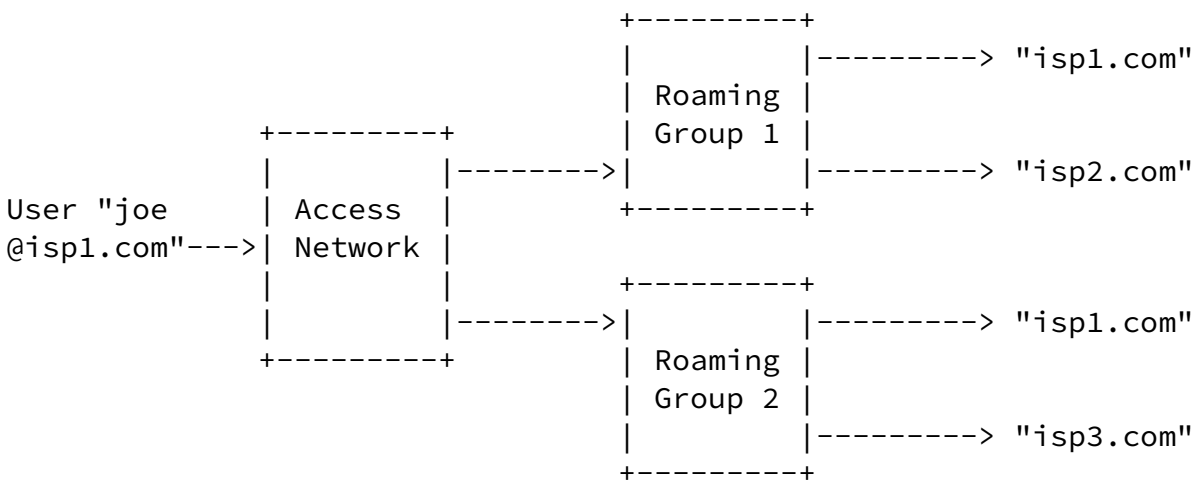

Figure 3: Ambiguous AAA routing

Currently planned networks include one level with a small number of intermediaries, just a few now and perhaps up to 50 as a maximum. However, multiple levels and higher number of alternative networks are possible in theory.

There has been requests to place credential and AAA route selection under user control, as the user is affected by the pricing and other differences. Optionally, automatic tools could make the selection based on the user's preferences. On the other hand, user control is similar to source routing, and as discussed earlier, network-based routing mechanisms have traditionally won over source routing-based mechanisms.

If users can control the selection of intermediaries, such intermediaries still have to be legitimate AAA proxies. That is, an access network should not send a request to an unknown intermediary. If it has a business relationship with three intermediaries int1.com, int2.com, and int3.com, it will route your request through one of them, even if you tried to request routing through mitm.org. Thus, NAI-based source routing is not source routing in the classic sense. It is merely suggesting preferences among already established routes. If the route does not already exist, or is not feasible, then NAI-based source routing cannot establish it.

An additional issue is that even if the intermediaries are legitimate, they could be switched. For instance, an access network

could advertise that it has a deal with cheapintermediary.net, and then switch the user's selection to priceyintermediary.com instead. To make this relevant, the pricing would have to be based on the intermediary. Even if it were possible to secure this selection, it would not be possible to guarantee that QoS or other properties

claimed by the network were indeed provided. As a result, it may be useful to think of the intermediary selection only as a hint.

Only a limited amount of information about AAA routes can be dynamically communicated. It is necessary to retrieve network and intermediary names, but quality of service or pricing information is clearly something that would need to be pre-provisioned, or perhaps just available via the web. Similarly, dynamic communication of network names can not be expected to provide all possible home network names, as their number can be quite large globally.

2.4 Payload Routing

The access network, mediating AAA infrastructure, and the home network may all have a desire to affect the kind of treatment payload packets get. This includes filtering, prioritization, routing paths, and mandatory tunneling.

Traditionally, the involved entities have been able to provide some control of this through AAA protocols such as RADIUS [6] and Diameter [9]. RFC 2868 [7] defines tunneling attributes that the home and mediating networks can use to establish mandatory tunneling at the access network. RFC 3588 [9] defines a filter syntax which can be used to install per-session filters under the control of AAA.

2.4.1 Issues with Payload Routing

In an attack described by Michael Richardson, a rogue hotspot operator (but with a legitimate roaming relationship to a home network) steals revenues from local hotspot operator by spoofing its identity. The rogue operator places a node with two interfaces in the coverage area of the local operator, and attaches to the Internet via this operator using a flat fee -based account. It then starts to advertise itself as a hotspot operator on the other interface, luring unsuspecting clients to use this node rather the than the local operator. The users authenticate via this node and its roaming

relationship. All AAA and payload traffic goes via the local hotspot, suitably NATted by the rogue node. As a result, the rogue operator gets the roaming service fees for a number of clients, whereas the local operator gets just one client.

Due to the way that the IEEE 802.11, IETF protocols, and common EAP methods have been designed, the rogue operator can actually advertise the same identity (SSID) as the local operator; the parameters advertised by the access point information are not authenticated end-to-end to the home network. EAP methods that support channel bindings [10] do not have this problem, but this support is not present in commonly used methods. Rogue access point can present a

different set of parameters to the client and to the home network. The current network access mechanisms enable us to have authentication, and link layer protection. They do not, however, guarantee anything about the delivery of the actual payload packets. In particular, there is no guarantee that the payload packets are delivered through a right route, or NATed only up to some specific number of times.

We call this the "payload route binding problem". In this problem, authentication and link layer support do not guarantee that the packets are actually routed through the path that appears to have been offered. Basically, if the "rogue" access point has a flat fee account and a contract with a clearing house, it can offer access to anyone with a per-byte account, NAT their packets, and send the traffic forward on its own account, and generate income. The local ISP would not be able to detect this by looking at the traffic stream. The user would not be able to detect it either, unless an EAP method with channel binding support is used. And even if it is, the user may not care about the identity of the access point enough to look at it; channel bindings can only ensure that all parties agree about the parameters, not that they make sense.

The working group did not come to a conclusion whether this attack needs to be prevented. Some of the proposed solutions include the use of EAP EMSK in the authentication exchange with the DHCP server, or the use of EAP to provide the certificates that SEND requires for the authentication of Router Advertisements. Either approach means that we are sure we are speaking at layer 3 to the services that we authenticated at layer two. However, this does not prevent an

attacker from offering such services, and then performing a NAT on the packets later. However, IPsec could be used to detect the presence of such NATs, even if NAT traversal is in use.

2.5 Discovery, Decision, and Selection

An alternative decomposition of the problem is to consider the discovery, decision, and selection aspects separately. Discovery consists of discovering access networks and associated POPs, discovering what identities the access networks will accept (either directly or through roaming relationships), and discovering which potential AAA intermediaries or routes exist.

Selection consists of attaching to the "right" access network and POP, offering an identity through EAP Identity Response, and providing a hint about the desired AAA intermediary. The selection of the AAA intermediary, along with the home and access networks, determines also the treatment of payload packets.

Decision can be either manual selection or automatic. Most likely, automatic mechanisms are preferred, even if manual selection should be retained as a fallback. The type of the decision also places additional requirements on the type of information that the discovery phase must provide. Just knowing which choices are available is probably enough for manual selection. Unfortunately, automatic selection based on a list of choices is by itself not possible:

- Some access networks may be preferred over others. For instance, the user's private corporate network may be preferred over a public network due to cost and efficiency reasons.
- Similarly, some credentials may be preferred over others.
- Use of an access network with direct connection to home network may be preferred over using mediating networks.
- Some mediating networks may be preferred to others, most likely based on cost. Note that optimizing cost is not a trivial problem, because the total cost may be a combination of a fixed fee, per-minute, per-megabyte, volume discounts, and so on.

- Preferences may come from the user, his or her employer (who's paying the bill), home network, or access network.

Different discovery mechanisms can accommodate such preferences in various ways. Some user input or perhaps a pre-provisioned database seems inevitable.

3. Design Constrains

All solutions to the network selection and discovery problem must satisfy the following additional constraints:

- AAA routing mechanisms have to work for requests, responses, as well as server-initiated messages.
- Solution is compatible with all AAA protocols.
- Does not prevent the introduction of new AAA or access network features, such as link-state AAA routing protocols or fast handoffs.
- Does not consume a significant amount of resources, such as

bandwidth or increase network attachment time.

- Does not cause a problem with limited packet sizes of current protocols.

- Where new protocol mechanisms are required, it should be possible to deploy the solution without requiring changes to the largest base of installed devices -- network access servers, wireless access points, and clients.

- Similarly, new solutions should allow interoperability with clients, access networks, AAA proxies, and AAA servers that have not been modified to support network discovery and selection.

4. Existing Work

4.1 IETF

There has already been a lot of past work in this area, including a number of IETF Proposed Standards generated by the ROAMOPS WG. The topic of roaming was considered different enough from both AAA and access protocols such as PPP that it deserved its own WG.

In addition to work on ROAMOPS directly relating to the problem, there has been work in SEAMOBY relating to scaling of network discovery mechanisms; work in PKIX relating to identity and credential selection; and work in AAA WG relating to access routing.

The PANA protocol [14] has a mechanism to advertise and select "ISPs" through the exchange of the ISP-Information AVP in its initial exchange.

Adrangi et al [16] discuss the use of the EAP-Request/Identity for network discovery. As noted in [10] Section 3.1, the minimum EAP MTU is 1020 octets, so that an EAP-Request/Identity is only guaranteed to be able to include 1015 octets within the Type-Data field. Since RFC 1035 [1] enables FQDNs to be up to 255 octets in length, this may not enable the announcement of very many networks, although if SSIDs are used, the maximum length of 32 octets per SSID may provide substantially better scaling. The use of other network identifiers than domain names is also possible, for instance the PANA protocol uses an a free form string and an SMI Network Management Private Enterprise Code [14], or Mobile Network Codes could be used as hinted in [16].

As noted in [30], the use of the EAP-Request/Identity for network discovery has substantial negative impact on handoff latency, since this may result in a station needing to initiate an EAP conversation with each Access Point in order to receive an EAP-Request/Identity describing which networks are supported. Since IEEE 802.11-1999 does not support use of Class 1 data frames in State 1 (unauthenticated, unassociated) within an ESS, this implies either that the APs must support 802.1X pre-authentication (optional in IEEE 802.11i) or that the station must associate with each AP prior to sending an EAPOL-Start to initiate EAP. This will dramatically increase handoff latency.

The effects to handoff latency depend also on the specific protocol design, and the expected likelihood of having to provide advertisements and initiate scanning of several access points. The use of advertisements only as a last resort when the AAA routing has failed is a better approach than the use of advertisement - scanning

procedure on every attachment.

Furthermore, if the AP has not been updated to present an up to date set of networks in the EAP-Requests/Identity, after associating to candidate APs and then choosing one, it is possible that the station will find that the chosen network is not supported after all. In this case, the station's EAP-Response/Identity may be answered with an updated EAP-Request/Identity, adding yet more latency.

4.2 IEEE

There has been work in IEEE 802.11 and 802.1 relating to network discovery enhancements.

Some recent contributions in this space include the following:

- [18] defines the Beacon and Probe Response mechanisms used with IEEE 802.11. Unfortunately, Beacons are only sent at the lowest supported rate. Studies such as [31] have identified MAC layer performance problems, and [28] have identified scaling issues resulting from a lowering of the Beacon interval.

- [21] discusses the evolution of authentication models in WLANs, and the need for the network to migrate from existing models to new ones, based on either EAP layer indications or through the use of SSIDs to represent more than the local network. It notes the potential need for management or structuring of the SSID space.

  The paper also notes that virtual APs have scalability issues. It does not analyze these scalability issues in relation to those existing in other alternative solutions, however.

- [22] discusses requirements for differentiation in the way that the user's payload traffic is routed, based on home network control. Such requirements have come up, for instance, in the context of 3GPP.

- [19] discusses mechanisms currently used to provide "Virtual AP" capabilities within a single physical access point. A "Virtual AP" appears at the MAC and IP layers to be distinct physical AP. As noted in the paper, full compatibility with existing 802.11 station implementations can only be maintained if each virtual AP uses a distinct MAC address (BSSID) for use in Beacons and Probe Responses. This draft does not discuss scaling issues in detail, but recommends that only a limited number of virtual APs be supported by a single physical access point. The simulations presented in [28] appear to confirm this conclusion; with a Beacon interval of 100 ms, once more than 8 virtual APs are supported on

a single channel, more than 20% of bandwidth is used for Beacons alone. This would indicate a limit of approximately 20 virtual APs per physical AP.

4.3 3GPP

The 3GPP technical specification [23] covers the interworking of WLAN networks with 2G and 3G networks. This specification discusses also network discovery and selection issues.

The specification requires that Access Network Discovery is performed as specified in the standards for the relevant WLAN link layer technology. An early version of the technical specification required the use of a 3GPP-specific SSID, but that has since then been abandoned; access network or local 3GPP network based SSIDs may be used instead. It has not been decided whether some conventions on the format of these SSIDs is required by 3GPP.

In addition to Access Network Discovery, it is necessary to select intermediary networks for the purposes of AAA Routing. In 3GPP, these networks are called Visited Public Land-based Mobile Networks (VPLMNs), and it is assumed that WLAN networks may have a contract with more than one VPLMN. GSM/UMTS roaming mechanisms are then employed for routing AAA requests from the VPLMN to the home network.

In order to select the VPLMN, the following is required:

- User can choose the desired VPLMN.
- AAA message are routed according to the NAI.
- Existing EAP mechanisms are used where possible.
- Extensibility is desired, to allow the advertisement of other parameters later.

The referenced 3GPP technical specification is a so called stage 2 specification, and contains only the principles of operation, leaving detailed protocol work for later. Nevertheless, the specification states that advertisement information shall be provided only when the access network is unable to route the request using normal AAA routing means, such as when it sees an unknown NAI domain. It is also stated that where VPLMN control is required, the necessary information is added to a NAI.

The security properties related to different mediating network selection mechanisms have been discussed in the 3GPP contribution

[24], which concludes that both SSID- and EAP-based mechanisms have roughly similar (and very limited) security properties, and that, as a result, network advertisement should be considered only as hints.

Ahmavaara, Haverinen, and Pichna [26] discuss the new network selection requirements that 3G-WLAN roaming introduces. It is necessary to support automatic network selection, and not just manual selection by the user. There may be multiple levels of networks, the hotspot owner may have a contract with a provider who in turn has a contract with one 3G network, and this 3G network has a roaming capability with a number of other networks.

4.4 Other

[27] discusses the need for network selection in a situation where there is more than one available access network with a roaming agreement to the home network. It also lists EAP-level, SSID-based, and PEAP-based mechanisms as potential solutions to the network selection problem.

Eijk et al [25] discussed the general issue of network selection. They concentrated primarily on the Access Network Discovery problem, based on various criteria, and did not consider the other aspects of the network selection problem. Nevertheless, they mention that one of the network selection problems is that the information about accessibility and roaming relationships is not stored in one location, but rather spread around the network.

5. Conclusions

The issues surrounding the network discovery and selection problem have been summarized.

In the opinion of the editors of this document, the main findings are:

- There is a clear need for access network discovery, identifier selection, AAA routing, and payload routing.

- Existing mechanisms appear sufficient for the control of payload routing, but there appears to be justification for enhanced mechanisms relating to access network discovery, identifier selection, and AAA routing.

- Nevertheless, many of the problems discussed in this draft are very hard when one considers them in an environment that requires a potentially large number of networks, fast handoffs, and automatic decisions.

- The proliferation of multiple competing network discovery technologies within IEEE 802, IETF, and 3GPP appears to a significant problem going forward. In the absence of a clearly defined solution to the problem it is likely that any or all of these solutions will be utilized, resulting in industry fragmentation and lack of interoperability.

  In order to avoid this fate, it is strongly suggested that a discussion be initiated between IETF and IEEE 802 in order to work out the roles of the each organization in solving this problem, and to invite 3GPP participation so that their requirements can be

fulfilled by the planned solutions.

- New link layers should be designed with facilities that enable the efficient distribution of network advertisement information.

- Solving all problems with current link layers and existing network access devices may not be possible. It may be useful to consider a phased approach where only certain functions are provided now, and the full functionality is provided when extensions to current link layers become available.

We will briefly comment on the specific mechanisms related to network discovery and selection:

- As noted in studies such as [31] and [28], the IEEE 802.11 Beacon/ Probe Response mechanism has substantial scaling issues, and as a

result a single physical access point is in practice limited to less than a dozen virtual APs on each channel with IEEE 802.11b.

The situation is improved substantially with successors such as IEEE 802.11a which enable additional channels, thus potentially increasing the number of potential virtual APs.

However, even these enhancements it is not feasible to advertise more than 50 different networks using existing mechanisms, and probably significantly less in most circumstances.

As a result, there appears to be justification for enhancing the scalability of network advertisements.

- Work is already underway in IEEE 802.1 to provide enhanced discovery functionality. For example, IEEE 802.1ab enables network devices to announce themselves and provide information on their capabilities. Similarly, the IEEE 802.1af has discussed the idea of supporting network discovery within a future revision to IEEE 802.1X. However, neither IEEE 801.ab nor IEEE 802.1af is likely to address the transport of large quantities of data where fragmentation would be a problem.

- Given that EAP does not support fragmentation of EAP-Request/ Identity packets, and that use of EAP for network selection on all

attachments will have a very substantial adverse impact on roaming performance without appropriate lower layer support (such as support for Class 1 data frames within IEEE 802.11), the use of EAP is at best limited. Long-term, it makes more sense for the desired functionality to be handled either within IEEE 802 or at the IP layer.

- In the IETF, a potential alternative is use of the SEAMOBY CARD protocol [13], which enables advertisement of network device capabilities over IP. Another alternative is the recently proposed Device Discovery Protocol (DDP) [12], which provides functionality equivalent to IEEE 802.1ab using ASN.1 encoded advertisements sent to a link-local scope multicast address.

  A limitation of these IP layer solutions is that they can only work as a means to speed up the attachment procedures when moving from one location to another; when a node starts up, it needs to be able to attach to a network before IP communications are available. This is fine for optimizations, but precludes the use in a case where the discovery information is mandatory before successful attachment can be accomplished, for instance when the access network is unable to route the AAA request unaided.

- "Phone-book" based approaches such as RFC 3017 appear attractive due to their ability to provide sufficient information for automatic selection decisions. However, there is no experience on applying such approaches to wireless access. The number of WLAN access points is significantly higher than the number of dial-in POPs; the distributed nature of the access network has created a more complicated business and roaming structure, and the expected rate of change in the information is high.

Finally, to address some of the security concerns that have come up during this work, the IETF should in any case initiate work that enables support for channel bindings in methods. Preferably, popular methods should be updated, ensuring compatibility with existing deployments. The representation of link layer parameters within EAP should utilize a common framework, to make it easier to define new link layers and keep the selection of EAP methods independent of the link layer.

6. Security Considerations

All aspects of the network discovery and selection problem are security related. The security issues and requirements have been discussed in the previous sections.

The security requirements for network discovery depend on the type of information being discovered. Some of the parameters may have a security impact, such as the claimed name of the network the user tries to attach to. Unfortunately, current EAP methods do not always make the verification of such parameters possible.

The security requirements for network selection depend on whether the selection is considered as a command or a hint. For instance, the selection that the user provided may be ignored if it relates to AAA routing and the access network can route the AAA traffic to the correct home network using other means in any case.

Normative References

[1] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[3] Aboba, B., Lu, J., Alsop, J., Ding, J. and W. Wang, "Review of Roaming Implementations", RFC 2194, September 1997.

[4] Aboba, B. and M. Beadles, "The Network Access Identifier", RFC 2486, January 1999.

[5] Aboba, B. and J. Vollbrecht, "Proxy Chaining and Policy Implementation in Roaming", RFC 2607, June 1999.

[6] Rigney, C., Willens, S., Rubens, A. and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[7] Zorn, G., Leifer, D., Rubens, A., Shriver, J., Holdrege, M. and I. Goyret, "RADIUS Attributes for Tunnel Protocol Support", RFC 2868, June 2000.

[8] Riegel, M. and G. Zorn, "XML DTD for Roaming Access Phone Book", RFC 3017, December 2000.

[9] Calhoun, P., Loughney, J., Guttman, E., Zorn, G. and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.

[10] Blunk, L., "Extensible Authentication Protocol (EAP)", draft-ietf-eap-rfc2284bis-07 (work in progress), December 2003.

[11] Housley, R. and T. Moore, "Certificate Extensions and Attributes Supporting Authentication in PPP and Wireless LAN", draft-ietf-pkix-wlan-extns-04 (work in progress), December 2002.

Informative References

[12] Enns, R., Marques, P. and D. Morrell, "Device Discovery Protocol (DDP)", draft-marques-ddp-00 (work in progress), May 2003.

[13] Liebsch, M., "Candidate Access Router Discovery", draft-ietf-seamoby-card-protocol-05 (work in progress), November 2003.

[14] Forsberg, D., "Protocol for Carrying Authentication for Network Access (PANA)", draft-ietf-pana-pana-02 (work in progress), October 2003.

[15] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", draft-ietf-ipsec-ikev2-07 (work in progress), April 2003.

[16] Adrangi, F., "Network Discovery and Selection within the EAP Framework", draft-adrangi-eap-network-discovery-and-selection-00 (work in progress), October 2003.

[17] Institute of Electrical and Electronics Engineers, "Local and Metropolitan Area Networks: Port-Based Network Access Control", IEEE Standard 802.1X, September 2001.

[18] Institute of Electrical and Electronics Engineers, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Standard 802.11, 1999.

[19] Aboba, B., "Virtual Access Points", IEEE Contribution 11-03-154r1, May 2003.

[20] Mishra, A., "Improving the latnecy of the Probe Phase during 802.11 Handoff", IEEE Contribution 11-03-417r2, May 2003.

[21] Hepworth, E., "Co-existence of Different Authentication Models", IEEE Contribution 11-03-0827 2003.

[22] Hong, C. and T. Yew, "Interworking - WLAN Control", IEEE Contribution 11-03-0843 2003.

[23] 3GPP, "3GPP System to Wireless Local Area Network (WLAN) interworking; System Description; Release 6", 3GPP Draft Technical Specification 23.234 v 2.2.0, December 2003.

[24] Ericsson, "Security of EAP and SSID based network

advertisements", 3GPP Contribution S3-030736, November 2003.

[25] Eijk, R., Brok, J., Bemmel, J. and B. Busropan, "Access Network Selection in a 4G Environment and the Role of Terminal and Service Platform", 10th WWRF, New York, October 2003.

[26] Ahmavaara, K., Haverinen, H. and R. Pichna, "Interworking Architecture between WLAN and 3G Systems", IEEE Communications Magazine, November 2003.

[27] Intel, "Wireless LAN (WLAN) End to End Guidelines for Enterprises and Public Hotspot Service Providers", November 2003.

[28] Velayos, H. and G. Karlsson, "Techniques to Reduce IEEE 802.11b MAC Layer Handover Time", Laboratory for Communication Networks, KTH, Royal Institute of Technology, Stockholm, Sweden, TRITA-IMIT-LCN R 03:02, April 2003.

[29] Judd, G. and P. Steenkiste, "Fixing 802.11 Access Point Selection", Sigcomm Poster Session 2002.

[30] Eronen, P., "Network Selection Issues", presentation to EAP WG at IETF 58, November 2003.

[31] Heusse, M., "Performance Anomaly of 802.11b", LSR-IMAG Laboratory, Grenoble, France, IEEE Infocom 2003.

Authors' Addresses

Jari Arkko
Ericsson

Jorvas 02420
Finland

EMail: jari.arkko@ericsson.com

Bernard Aboba
Microsoft

One Microsoft Way
Redmond, WA 98052
USA

EMail: aboba@internaut.com

Appendix A. Contributors

This draft is based on the discussion held on the EAP WG mailing list in December 2003, and on a number of input documents such as [16]. The editors of this document would like to especially acknowledge the contributions of Farid Adrangi, Farooq Bari, Michael Richardson, Pasi Eronen, Mark Watson, Mark Grayson, Johan Rune, and Tomas Goldbeck-Lowe.

Intellectual Property Statement

Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.