

Network Working Group
Request for Comments: 5113
Category: Informational

J. Arkko
Ericsson
B. Aboba
Microsoft
J. Korhonen, Ed.
TeliaSonera
F. Bari
AT&T
January 2008

Network Discovery and Selection Problem

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

When multiple access networks are available, users may have difficulty in selecting which network to connect to and how to authenticate with that network. This document defines the network discovery and selection problem, dividing it into multiple sub-problems. Some constraints on potential solutions are outlined, and the limitations of several solutions (including existing ones) are discussed.

Table of Contents

1. Introduction	3
1.1. Terminology Used in This Document	4
2. Problem Definition	7
2.1. Discovery of Points of Attachment	8
2.2. Identity Selection	9
2.3. AAA Routing	11
2.3.1. The Default Free Zone	13
2.3.2. Route Selection and Policy	14
2.3.3. Source Routing	15
2.4. Network Capabilities Discovery	17
3. Design Issues	18
3.1. AAA Routing	18
3.2. Backward Compatibility	18
3.3. Efficiency Constraints	19
3.4. Scalability	19
3.5. Static Versus Dynamic Discovery	21
3.6. Security	21
3.7. Management	22
4. Conclusions	23
5. Security Considerations	25
6. Informative References	26
Appendix A. Existing Work	32
A.1. IETF	32
A.2. IEEE 802	33
A.3. 3GPP	35
A.4. Other	36
Appendix B. Acknowledgements	37

1. Introduction

Today, network access clients are typically pre-configured with a list of access networks and corresponding identities and credentials. However, as network access mechanisms and operators have proliferated, it has become increasingly likely that users will encounter networks for which no pre-configured settings are available, yet which offer desired services and the ability to successfully authenticate with the user's home realm. It is also possible that pre-configured settings will not be adequate in some situations. In such a situation, users can have difficulty in determining which network to connect to, and how to authenticate to that network.

The problem arises when any of the following conditions are true:

- o Within a single network, more than one network attachment point is available, and the attachment points differ in their roaming arrangements, or access to services. While the link layer capabilities of a point of attachment may be advertised, higher-layer capabilities, such as roaming arrangements, end-to-end quality of service, or Internet access restrictions, may not be. As a result, a user may have difficulty determining which services are available at each network attachment point, and which attachment points it can successfully authenticate to. For example, it is possible that a roaming agreement will only enable a user to authenticate to the home realm from some points of attachment, but not others. Similarly, it is possible that access to the Internet may be restricted at some points of attachment, but not others, or that end-to-end quality of service may not be available in all locations. In these situations, the network access client cannot assume that all points of attachment within a network offer identical capabilities.
- o Multiple networks are available for which the user has no corresponding pre-configuration. The user may not have pre-configured an identity and associated credentials for use with a network, yet it is possible that the user's home realm is reachable from that network, enabling the user to successfully authenticate. However, unless the roaming arrangements are advertised, the network access client cannot determine a priori whether successful authentication is likely. In this situation, it is possible that the user will need to try multiple networks in order to find one to which it can successfully authenticate, or it is possible that the user will not be able to obtain access at all, even though successful authentication is feasible.

- o The user has multiple sets of credentials. Where no pre-configuration exists, it is possible that the user will not be able to determine which credentials to use with which attachment point, or even whether any credentials it possesses will allow it to authenticate successfully. An identity and associated credentials can be usable for authentication with multiple networks, and not all of these networks will be pre-configured. For example, the user could have one set of credentials from a public service provider and another set from an employer, and a network might enable authentication with one or more of these credentials. Yet, without pre-configuration, multiple unsuccessful authentication attempts could be needed for each attachment point in order to determine what credentials are usable, wasting valuable time and resulting in user frustration. In order to choose between multiple attachment points, it can be helpful to provide additional information to enable the correct credentials to be determined.
- o There are multiple potential roaming paths between the visited realm and the user's home realm, and service parameters or pricing differs between them. In this situation, there could be multiple ways for the user to successfully authenticate using the same identity and credentials, yet the cost of each approach might differ. In this case, the access network may not be able to determine the roaming path that best matches the user's preferences. This can lead to the user being charged more than necessary, or not obtaining the desired services. For example, the visited access realm could have both a direct relationship with the home realm and an indirect relationship through a roaming consortium. Current Authentication, Authorization, and Accounting (AAA) protocols may not be able to route the access request to the home AAA server purely based on the realm within the Network Access Identifier (NAI) [RFC4282]. In addition, payload packets can be routed or tunneled differently, based on the roaming relationship path. This may have an impact on the available services or their pricing.

In Section 2, the network discovery and selection problem is defined and divided into sub-problems. Some solution constraints are outlined in Section 3. Section 4 provides conclusions and suggestions for future work. Appendix A discusses existing solutions to portions of the problem.

1.1. Terminology Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Authentication, Authorization, and Accounting (AAA)

AAA protocols with EAP support include Remote Authentication Dial-In User Service (RADIUS) [RFC3579] and Diameter [RFC4072].

Access Point (AP)

An entity that has station functionality and provides access to distribution services via the wireless medium (WM) for associated stations.

Access Technology Selection

This refers to the selection between access technologies, e.g., 802.11, Universal Mobile Telecommunications System (UMTS), WiMAX. The selection will be dependent upon the access technologies supported by the device and the availability of networks supporting those technologies.

Bearer Selection

For some access technologies (e.g., UMTS), there can be a possibility for delivery of a service (e.g., voice) by using either a circuit-switched or packet-switched bearer. Bearer selection refers to selecting one of the bearer types for service delivery. The decision can be based on support of the bearer type by the device and the network as well as user subscription and operator preferences.

Basic Service Set (BSS)

A set of stations controlled by a single coordination function.

Decorated NAI

A NAI specifying a source route. See Section 2.7 of RFC 4282 [RFC4282] for more information.

Extended Service Set (ESS)

A set of one or more interconnected basic service sets (BSSs) with the same Service Set Identifier (SSID) and integrated local area networks (LANs), which appears as a single BSS to the logical link control layer at any station associated with one of those BSSs. This refers to a mechanism that a node uses to discover the networks that are reachable from a given access network.

Network Access Identifier (NAI)

The Network Access Identifier (NAI) [RFC4282] is the user identity submitted by the client during network access authentication. In roaming, the purpose of the NAI is to identify the user as well as to assist in the routing of the authentication request. Please note that the NAI may not necessarily be the same as the user's e-mail address or the user identity submitted in an application layer authentication.

Network Access Server (NAS)

The device that peers connect to in order to obtain access to the network. In Point-to-Point Tunneling Protocol (PPTP) terminology, this is referred to as the PPTP Access Concentrator (PAC), and in Layer 2 Tunneling Protocol (L2TP) terminology, it is referred to as the L2TP Access Concentrator (LAC). In IEEE 802.11, it is referred to as an Access Point (AP).

Network Discovery

The mechanism used to discover available networks. The discovery mechanism may be passive or active, and depends on the access technology. In passive network discovery, the node listens for network announcements; in active network discovery, the node solicits network announcements. It is possible for an access technology to utilize both passive and active network discovery mechanisms.

Network Selection

Selection of an operator/ISP for network access. Network selection occurs prior to network access authentication.

Realm

The realm portion of an NAI [RFC4282].

Realm Selection

The selection of the realm (and corresponding NAI) used to access the network. A realm can be reachable from more than one access network type, and selection of a realm may not enable network capabilities.

Roaming Capability

Roaming capability can be loosely defined as the ability to use any one of multiple Internet Service Providers (ISPs), while maintaining a formal, customer-vendor relationship with only one. Examples of cases where roaming capability might be required include ISP "confederations" and ISP-provided corporate network access support.

Station (STA)

A device that contains an IEEE 802.11 conformant medium access control (MAC) and physical layer (PHY) interface to the wireless medium (WM).

2. Problem Definition

The network discovery and selection problem can be broken down into multiple sub-problems. These include:

- o Discovery of points of attachment. This involves the discovery of points of attachment in the vicinity, as well as their capabilities.
- o Identifier selection. This involves selection of the NAI (and credentials) used to authenticate to the selected point of attachment.
- o AAA routing. This involves routing of the AAA conversation back to the home AAA server, based on the realm of the selected NAI.
- o Payload routing. This involves the routing of data packets, in the situation where mechanisms more advanced than destination-based routing are required. While this problem is interesting, it is not discussed further in this document.
- o Network capability discovery. This involves discovering the capabilities of an access network, such as whether certain services are reachable through the access network and the charging policy.

Alternatively, the problem can be divided into discovery, decision, and selection components. The AAA routing problem, for instance, involves all components: discovery (which mediating networks are available), decision (choosing the "best" one), and selection (selecting which mediating network to use) components.

2.1. Discovery of Points of Attachment

Traditionally, the discovery of points of attachment has been handled by out-of-band mechanisms or link or network layer advertisements.

RFC 2194 [RFC2194] describes the pre-provisioning of dial-up 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 [RFC3017] describes the IETF Proposed Standard for the Roaming Access eXtensible Markup Language (XML) Document Type Definition (DTD). This covers not only the attributes of the Points of Presence (PoP) and Internet Service Providers (ISPs), but also hints on the appropriate NAI to be used with a particular PoP. The XML DTD supports dial-in and X.25 access, but has extensible address and media type fields.

As access networks and the points of attachment have proliferated, out-of-band pre-configuration has become increasingly difficult. For networks with many points of attachment, keeping a complete and up-to-date list of points of attachment can be difficult. As a result, wireless network access clients typically only attempt to pre-configure information relating to access networks, rather than individual points of attachment.

In IEEE 802.11 Wireless Local Area Networks (WLAN), the Beacon and Probe Request/Response mechanism provides a way for Stations to discover Access Points (AP) and the capabilities of those APs. The IEEE 802.11 specification [IEEE.802.11-2003] provides support for both passive (Beacon) and active (Probe Request/Response) discovery mechanisms; [Fixingapsell] studied the effectiveness of these mechanisms.

Among the Information Elements (IE) included within the Beacon and Probe Response is the Service Set Identifier (SSID), a non-unique identifier of the network to which an AP is attached. The Beacon/Probe facility therefore enables network discovery, as well as the discovery of points of attachment and the capabilities of those points of attachment.

The Global System for Mobile Communications (GSM) specifications also provide for discovery of points of attachment, as does the Candidate Access Router Discovery (CARD) [RFC4066] protocol developed by the IETF SEAMOBY Working Group (WG).

Along with discovery of points of attachment, the capabilities of access networks are also typically discovered. These may include:

- o Access network name (e.g., IEEE 802.11 SSID)
- o Lower layer security mechanism (e.g., IEEE 802.11 Wired Equivalent Privacy (WEP) vs. Wi-Fi Protected Access 2 (WPA2))
- o Quality of service capabilities (e.g., IEEE 802.11e support)
- o Bearer capabilities (e.g., circuit-switched, packet-switched, or both)

Even though pre-configuration of access networks scales better than pre-configuration of points of attachment, where many access networks can be used to authenticate to a home realm, providing complete and up-to-date information on each access network can be challenging.

In such a situation, network access client configuration can be minimized by providing information relating to each home realm, rather than each access network. One way to enable this is for an access network to support "virtual Access Points" (virtual APs), and for each point of attachment to support virtual APs corresponding to each reachable home realm.

While a single IEEE 802.11 network may only utilize a single SSID, it may cover a wide geographical area, and as a result, may be segmented, utilizing multiple prefixes. It is possible that a single SSID may be advertised on multiple channels, and may support multiple access mechanisms (including Universal Access Method (UAM) and IEEE 802.1X [IEEE.8021X-2004]) which may differ between points of attachment. A single SSID may also support dynamic VLAN access as described in [RFC3580], or may support authentication to multiple home AAA servers supporting different realms. As a result, users of a single point of attachment, connecting to the same SSID, may not have the same set of services available.

2.2. Identity Selection

As networks proliferate, it becomes more and more likely that a user may have multiple identities and credential sets, available for use in different situations. For example, the user may have an account with one or more Public WLAN providers, a corporate WLAN, and one or more wireless Wide Area Network (WAN) providers.

Typically, the user will choose an identity and corresponding credential set based on the selected network, perhaps with additional assistance provided by the chosen authentication mechanism. For example, if Extensible Authentication Protocol - Transport Layer Security (EAP-TLS) is the authentication mechanism used with a particular network, then the user will select the appropriate EAP-TLS

client certificate based, in part, on the list of trust anchors provided by the EAP-TLS server.

However, in access networks where roaming is enabled, the mapping between an access network and an identity/credential set may not be one to one. For example, it is possible for multiple identities to be usable on an access network, or for a given identity to be usable on a single access network, which may or may not be available.

Figure 1 illustrates a situation where a user identity may not be usable on a potential access network. In this case, access network 1 enables access to users within the realm "isp1.example.com", whereas access network 3 enables access to users within the realm "corp2.example.com"; access network 2 enables access to users within both realms.




Figure 1: Two credentials, three possible access networks

In this situation, a user only possessing an identity within the "corp2.example.com" realm can only successfully authenticate to access networks 2 or 3; a user possessing an identity within the "isp1.example.com" realm can only successfully authenticate to access networks 1 or 2; a user possessing identities within both realms can connect to any of the access networks. The question is: how does the user figure out which access networks it can successfully authenticate to, preferably prior to choosing a point of attachment?

Traditionally, hints useful in identity selection have been provided out-of-band. For example, the XML DTD, described in [RFC3017], enables a client to select between potential points of attachment as

well as to select the NAI and credentials to use in authenticating with it.

Where all points of attachment within an access network enable authentication utilizing a set of realms, selection of an access network provides knowledge of the identities that a client can use to successfully authenticate. For example, in an access network, the set of supported realms corresponding to network name can be pre-configured.

In some cases, it may not be possible to determine the available access networks prior to authentication. For example, [IEEE.8021X-2004] does not support network discovery on IEEE 802 wired networks, so that the peer cannot determine which access network it has connected to prior to the initiation of the EAP exchange.

It is also possible for hints to be embedded within credentials. In [RFC4334], 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.

However, there may be situations in which an access network may not accept a static set of realms at every point of attachment. For example, as part of a roaming agreement, only points of attachment within a given region or country may be made available. In these situations, mechanisms such as hints embedded within credentials or pre-configuration of access network to realm mappings may not be sufficient. Instead, it is necessary for the client to discover usable identities dynamically.

This is the problem that RFC 4284 [RFC4284] attempts to solve, using the EAP-Request/Identity to communicate a list of supported realms. However, the problems inherent in this approach are many, as discussed in Appendix A.1.

Note that identity selection also implies selection of different credentials, and potentially, selection of different EAP authentication methods. In some situations this may imply serious security vulnerabilities. These are discussed in depth in Section 5.

2.3. AAA Routing

Once the identity has been selected, the AAA infrastructure needs to route the access request back to the home AAA server. Typically, the routing is based on the Network Access Identifier (NAI) defined in [RFC4282].

Where the NAI does not encode a source route, the routing of requests is determined by the AAA infrastructure. As described in [RFC2194], most roaming implementations are relatively simple, relying on a static realm routing table that determines the next hop based on the NAI realm included in the User-Name attribute within the Access-Request. Within RADIUS, the IP address of the home AAA server is typically determined based on static mappings of realms to IP addresses maintained within RADIUS proxies.

Diameter [RFC3588] supports mechanisms for intra- and inter-domain service discovery, including support for DNS as well as service discovery protocols such as Service Location Protocol version 2 (SLPv2) [RFC2608]. As a result, it may not be necessary to configure static tables mapping realms to the IP addresses of Diameter agents. However, while this simplifies maintenance of the AAA routing infrastructure, it does not necessarily simplify roaming-relationship path selection.

As noted in RFC 2607 [RFC2607], 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 between each AAA client and server.

Diameter [RFC3588] supports certificate-based authentication (using either TLS or IPsec) as well as Redirect functionality, enabling a Diameter client to obtain a referral to the home server from a Diameter redirect server, so that the client can contact the home server directly. In situations in which a trust model can be established, these Diameter capabilities can enable a reduction in the length of the roaming relationship path.

However, in practice there are a number of pitfalls. In order for certificate-based authentication to enable communication between a Network Access Server (NAS) or local proxy and the home AAA server, trust anchors need to be configured, and certificates need to be selected. The AAA server certificate needs to chain to a trust anchor configured on the AAA client, and the AAA client certificate needs to chain to a trust anchor configured on the AAA server. Where multiple potential roaming relationship paths are available, this will reflect itself in multiple certificate choices, transforming the path selection problem into a certificate selection problem. Depending on the functionality supported within the certificate selection implementation, this may not make the problem easier to solve. For example, in order to provide the desired control over the roaming path, it may be necessary to implement custom certificate selection logic, which may be difficult to introduce within a

certificate handling implementation designed for general-purpose usage.

As noted in [RFC4284], it is also possible to utilize an NAI for the purposes of source routing. In this case, the client provides guidance to the AAA infrastructure as to how it would like the access request to be routed. An NAI including source-routing information is said to be "decorated"; the decoration format is defined in [RFC4282].

When decoration is utilized, the EAP peer provides the decorated NAI within the EAP-Response/Identity, and as described in [RFC3579], the NAS copies the decorated NAI included in the EAP-Response/Identity into the User-Name attribute included within the access request. As the access request transits the roaming relationship path, AAA proxies determine the next hop based on the realm included within the User-Name attribute, in the process, successively removing decoration from the NAI included in the User-Name attribute. In contrast, the decorated NAI included within the EAP-Response/Identity encapsulated in the access request remains untouched. As a result, when the access request arrives at the AAA home server, the decorated NAI included in the EAP-Response/Identity may differ from the NAI included in the User-Name attribute (which may have some or all of the decoration removed). For the purpose of identity verification, the EAP server utilizes the NAI in the User-Name attribute, rather than the NAI in the EAP-Response/Identity.

Over the long term, it is expected that the need for NAI "decoration" and source routing will disappear. 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 Unix-to-Unix CoPy Protocol (UUCP) and FidoNet. Prior to the implementation of email gateways utilizing MX RR routing, email address-based source-routing was used extensively. However, over time the need for email source-routing disappeared.

2.3.1. The Default Free Zone

AAA clients on the edge of the network, such as NAS devices and local AAA proxies, typically maintain a default realm route, providing a default next hop for realms not otherwise taken into account within the realm routing table. This permits devices with limited resources to maintain a small realm routing table. Deeper within the AAA infrastructure, AAA proxies may be maintained with a "default free" realm table, listing next hops for all known realms, but not providing a default realm route.

While dynamic realm routing protocols are not in use within AAA infrastructure today, even if such protocols were to be introduced, it is likely that they would be deployed solely within the core AAA infrastructure, but not on NAS devices, which are typically resource constrained.

Since NAS devices do not maintain a full realm routing table, they do not have knowledge of all the realms reachable from the local network. The situation is analogous to that of Internet hosts or edge routers that do not participate in the BGP mesh. In order for an Internet host to determine whether it can reach a destination on the Internet, it is necessary to send a packet to the destination.

Similarly, when a user provides an NAI to the NAS, the NAS does not know a priori whether or not the realm encoded in the NAI is reachable; it simply forwards the access request to the next hop on the roaming relationship path. Eventually, the access request reaches the "default free" zone, where a core AAA proxy determines whether or not the realm is reachable. As described in [RFC4284], where EAP authentication is in use, the core AAA proxy can send an Access-Reject, or it can send an Access-Challenge encapsulating an EAP-Request/Identity containing "realm hints" based on the content of the "default free" realm routing table.

There are a number of intrinsic problems with this approach. Where the "default free" routing table is large, it may not fit within a single EAP packet, and the core AAA proxy may not have a mechanism for selecting the most promising entries to include. Even where the "default free" realm routing table would fit within a single EAP-Request/Identity packet, the core AAA router may not choose to include all entries, since the list of realm routes could be considered confidential information not appropriate for disclosure to hosts seeking network access. Therefore, it cannot be assumed that the list of "realm hints" included within the EAP-Request/Identity is complete. Given this, a NAS or local AAA proxy snooping the EAP-Request/Identity cannot rely on it to provide a complete list of reachable realms. The "realm hint" mechanism described in [RFC4284] is not a dynamic routing protocol.

2.3.2. Route Selection and Policy

Along with lack of a dynamic AAA routing protocol, today's AAA infrastructure lacks mechanisms for route selection and policy. As a result, multiple routes may exist to a destination realm, without a mechanism for the selection of a preferred route.

In Figure 2, Roaming Groups 1 and 2 both include a route to the realm "a.example.com". However, these realm routes are not disseminated to the NAS along with associated metrics, and, as a result, there is no mechanism for implementation of dynamic routing policies (such as selection of realm routes by shortest path, or preference for routes originating at a given proxy).




Figure 2: Multiple routes to a destination realm

In the example in Figure 2, access through Roaming Group 1 may be less expensive than access through Roaming Group 2, and as a result it would be desirable to prefer Roaming Group 1 as a next hop for an NAI with a realm of "a.example.com". However, the only way to obtain this result would be to manually configure the NAS realm routing table with the following entries:

Realm	Next Hop
-----	-----
b.example.com	Roaming Group 1
c.example.com	Roaming Group 2
Default	Roaming Group 1

While manual configuration may be practical in situations where the realm routing table is small and entries are static, where the list of supported realms change frequently, or the preferences change dynamically, manual configuration will not be manageable.

2.3.3. Source Routing

Due to the limitations of current AAA routing mechanisms, there are situations in which NAI-based source routing is used to influence the roaming relationship path. However, since the AAA proxies on the roaming relationship path are constrained by existing relationships, NAI-based source routing is not source routing in the classic sense;

it merely suggests preferences that the AAA proxy can choose not to accommodate.

Where realm routes are set up as the result of pre-configuration and dynamic route establishment is not supported, if a realm route does not exist, then NAI-based source routing cannot establish it. Even where dynamic route establishment is possible, such as where the AAA client and server support certificate-based authentication, and AAA servers are discoverable (such as via the mechanisms described in [RFC3588]), an AAA proxy may choose not to establish a realm route by initiating the discovery process based on a suggestion in an NAI-based source route.

Where the realm route does exist, or the AAA proxy is capable of establishing it dynamically, the AAA proxy may choose not to authorize the client to use it.

While, in principle, source routing can provide users with better control over AAA routing decisions, there are a number of practical problems to be overcome. In order to enable the client to construct optimal source routes, it is necessary for it to be provided with a complete and up-to-date realm routing table. However, if a solution to this problem were readily available, then it could be applied to the AAA routing infrastructure, enabling the selection of routes without the need for user intervention.

As noted in [Eronen04], only a limited number of parameters can be updated dynamically. For example, quality of service or pricing information typically will be pre-provisioned or made available on the web rather than being updated on a continuous basis. Where realm names are communicated dynamically, the "default free" realm list is unlikely to be provided in full since this table could be quite large. Given the constraints on the availability of information, the construction of source routes typically needs to occur in the face of incomplete knowledge.

In addition, there are few mechanisms available to audit whether the requested source route is honored by the AAA infrastructure. For example, an access network could advertise a realm route to "costsless.example.com", while instead routing the access-request through "costsmore.example.com". While the decorated NAI would be made available to the home AAA server in the EAP-Response/Identity, the home AAA server might have a difficult time verifying that the source route requested in the decorated NAI was actually honored by the AAA infrastructure. Similarly, it could be difficult to determine whether quality of service (QoS) or other routing requests were actually provided as requested. To some extent, this problem

may be addressed as part of the business arrangements between roaming partners, which may provide minimum service-level guarantees.

Given the potential issues with source routing, conventional AAA routing mechanisms are to be preferred wherever possible. Where an error is encountered, such as an attempt to authenticate to an unreachable realm, "realm hints" can be provided as described [RFC4284]. However, this approach has severe scalability limitations, as outlined in Appendix A.1.

2.4. Network Capabilities Discovery

Network capability discovery focuses on discovery of the services offered by networks, not just the capabilities of individual points of attachment. By acquiring additional information on access network characteristics, it is possible for users to make a more informed access decision. These characteristics may include:

- o Roaming relationships between the access network provider and other network providers and associated costs. Where the network access client is not pre-configured with an identity and credentials corresponding to a local access network, it will need to be able to determine whether one or more home realms are reachable from an access network so that successful authentication can be possible.
- o EAP authentication methods. While the EAP authentication methods supported by a home realm can only be determined by contacting the home AAA server, it is possible that the local realm will also support one or more EAP methods. For example, a user may be able to utilize EAP-SIM (Extensible Authentication Protocol - Subscriber Identity Module) to authenticate to the access network directly, rather than having to authenticate to the home network.
- o End-to-end quality of service capability. While local quality of service capabilities are typically advertised by the access network (e.g., support for Wi-Fi Multimedia (WMM)), the availability of end-to-end QoS services may not be advertised.
- o Service parameters, such as the existence of middleboxes or firewalls. If the network access client is not made aware of the Internet access that it will receive on connecting to a point of attachment, it is possible that the user may not be able to access the desired services.

Reference [IEEE.11-04-0624] classifies the possible steps at which IEEE 802.11 networks can acquire this information:

- o Pre-association
- o Post-association (or pre-authentication)
- o Post-authentication

In the interest of minimizing connectivity delays, all of the information required for network selection (including both access network capabilities and global characteristics) needs to be provided prior to authentication.

By the time authentication occurs, the node has typically selected the access network, the NAI to be used to authenticate, as well as the point of attachment. Should it learn information during the authentication process that would cause it to revise one or more of those decisions, the node will need to select a new network, point of attachment, and/or identity, and then go through the authentication process all over again. Such a process is likely to be both time consuming and unreliable.

3. Design Issues

The following factors should be taken into consideration while evaluating solutions to the problem of network selection and discovery.

3.1. AAA Routing

Solutions to the AAA routing issues discussed in Section 2.3 need to apply to a wide range of AAA messages, and should not restrict the introduction of new AAA or access network functionality. For example, AAA routing mechanisms should work for access requests and responses as well as accounting requests and responses and server-initiated messages. Solutions should not restrict the development of new AAA attributes, access types, or performance optimizations (such as fast handoff support).

3.2. Backward Compatibility

Solutions need to maintain backward compatibility. In particular:

- o Selection-aware clients need to interoperate with legacy NAS devices and AAA servers.
- o Selection-aware AAA infrastructure needs to interoperate with legacy clients and NAS devices.

For example, selection-aware clients should not transmit packets larger than legacy NAS devices or AAA servers can handle. Where protocol extensions are required, changes should be required to as few infrastructure elements as possible. For example, extensions that require upgrades to existing NAS devices will be more difficult to deploy than proposals that are incrementally deployable based on phased upgrades of clients or AAA servers.

3.3. Efficiency Constraints

Solutions should be efficient as measured by channel utilization, bandwidth consumption, handoff delay, and energy utilization. Mechanisms that depend on multicast frames need to be designed with care since multicast frames are often sent at the lowest supported rate and therefore consume considerable channel time as well as energy on the part of listening nodes. Depending on the deployment, it is possible for bandwidth to be constrained both on the link, as well as in the backend AAA infrastructure. As a result, chatty mechanisms such as keepalives or periodic probe packets are to be avoided. Given the volume handled by AAA servers, solutions should also be conscious of adding to the load, particularly in cases where this could enable denial-of-service attacks. For example, it would be a bad idea for a NAS to attempt to obtain an updated realm routing table by periodically sending probe EAP-Response/Identity packets to the AAA infrastructure in order to obtain "realm hints" as described in [RFC4284]. Not only would this add significant load to the AAA infrastructure (particularly in cases where the AAA server was already overloaded, thereby dropping packets resulting in retransmission by the NAS), but it would also not provide the NAS with a complete realm routing table, for reasons described in Section 2.3.

Battery consumption is a significant constraint for handheld devices. Therefore, mechanisms that require significant increases in packets transmitted, or the fraction of time during which the host needs to listen (such as proposals that require continuous scanning), are to be discouraged. In addition, the solution should not significantly impact the time required to complete network attachment.

3.4. Scalability

Given limitations on frame sizes and channel utilization, it is important that solutions scale less than linearly in terms of the number of networks and realms supported. For example, solutions such as [RFC4284] increase the size of advertisements in proportion to the number of entries in the realm routing table. This approach does not scale to support a large number of networks and realms.

Similarly, approaches that utilize separate Beacons for each "virtual AP" introduce additional Beacons in proportion to the number of networks being advertised. While such an approach may minimize the pre-configuration required for network access clients, the proliferation of "virtual APs" can result in high utilization of the wireless medium. For example, the 802.11 Beacon is sent only at a rate within the basic rate set, which typically consists of the lowest supported rates, or perhaps only the lowest supported rate. As a result, "virtual AP" mechanisms that require a separate Beacon for each "virtual AP" do not scale well.

For example, with a Beacon interval of 100 Time Units (TUs) or 102.4 ms (9.8 Beacons/second), twenty 802.11b "virtual APs" each announcing their own Beacon of 170 octets would result in a channel utilization of 37.9 percent. The calculation can be verified as follows:

1. A single 170-octet Beacon sent at 1 Mbps will utilize the channel for 1360 us (1360 bits @ 1 Mbps);
2. Adding 144 us for the Physical Layer Convergence Procedure (PLCP) long preamble (144 bits @ 1 Mbps), 48 us for the PLCP header (48 bits @ 1 Mbps), 10 us for the Short Interframe Space (SIFS), 50 us for the Distributed Interframe Space (DIFS), and 320 us for the average minimum Contention Window without backoff ($CW_{min}/2 * aSlotTime = 32/2 * 20$ us) implies that a single Beacon will utilize an 802.11b channel for 1932 us;
3. Multiply the channel time per Beacon by 196 Beacons/second, and we obtain a channel utilization of 378672 us/second = 37.9 percent.

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. Another issue with the Beacon and Probe Request/Response mechanism is that it is either insecure or its security can be assured only as part of authenticating to the network (e.g., verifying the advertised capabilities within the 4-way handshake).

A number of enhancements have been proposed to the Beacon/Probe Response mechanism in order to improve scalability and performance in roaming scenarios. These include allowing APs to announce capabilities of neighbor APs as well as their own [IEEE.802.11k]. More scalable mechanisms for support of "virtual APs" within IEEE 802.11 have also been proposed [IEEE.802.11v]; generally these proposals collapse multiple "virtual AP" advertisements into a single advertisement.

Higher-layer mechanisms can also be used to improve scalability since, by running over IP, they can utilize facilities, such as fragmentation, that may not be available at the link layer. For example, in IEEE 802.11, Beacon frames cannot use fragmentation because they are multicast frames.

3.5. Static Versus Dynamic Discovery

"Phone-book" based approaches such as [RFC3017] can provide information for automatic selection decisions. While this approach has been applied to wireless access, it typically can only be used successfully within a single operator or limited roaming partner deployment. For example, were a "Phone-Book" approach to attempt to incorporate information from a large number of roaming partners, it could become quite difficult to keep the information simultaneously comprehensive and up to date. As noted in [Priest04] and [GROETING], a large fraction of current WLAN access points operate on the default SSID, which may make it difficult to distinguish roaming partner networks by SSID. In any case, in wireless networks, dynamic discovery is a practical requirement since a node needs to know which APs are within range before it can connect.

3.6. Security

Network discovery and selection mechanisms may introduce new security vulnerabilities. As noted in Section 2.3.1, network operators may consider the AAA routing table to be confidential information, and therefore may not wish to provide it to unauthenticated peers via the mechanism described in RFC 4284. While the peer could provide a list of the realms it supports, with the authenticator choosing one, this approach raises privacy concerns. Since identity selection occurs prior to authentication, the peer's supported realms would be sent in cleartext, enabling an attacker to determine the realms for which a potential victim has credentials. This risk can be mitigated by restricting peer disclosure. For example, a peer may only disclose additional realms in situations where an initially selected identity has proved unusable.

Since network selection occurs prior to authentication, it is typically not possible to secure mechanisms for network discovery or identity selection, although it may be possible to provide for secure confirmation after authentication is complete. As an example, some parameters discovered during network discovery may be confirmable via EAP Channel Bindings; others may be confirmed in a subsequent Secure Association Protocol handshake.

However, there are situations in which advertised parameters may not be confirmable. This could lead to "bidding down" vulnerabilities. Section 7.8 of [RFC3748] states:

Within or associated with each authenticator, it is not anticipated that a particular named peer will support a choice of methods. This would make the peer vulnerable to attacks that negotiate the least secure method from among a set. Instead, for each named peer, there SHOULD be an indication of exactly one method used to authenticate that peer name. If a peer needs to make use of different authentication methods under different circumstances, then distinct identities SHOULD be employed, each of which identifies exactly one authentication method.

In practice, where the authenticator operates in "pass-through" mode, the EAP method negotiation will occur between the EAP peer and server, and therefore the peer will need to associate a single EAP method with a given EAP server. Where multiple EAP servers and corresponding identities may be reachable from the same selected network, the EAP peer may have difficulty determining which identity (and corresponding EAP method) should be used. Unlike network selection, which may be securely confirmed within a Secure Association Protocol handshake, identity selection hints provided within the EAP-Request/Identity are not secured.

As a result, where the identity selection mechanism described in RFC 4284 is used, the "hints" provided could be used by an attacker to convince the victim to select an identity corresponding to an EAP method offering lesser security (e.g., EAP MD5-Challenge). One way to mitigate this risk is for the peer to only utilize EAP methods satisfying the [RFC4017] security requirements, and for the peer to select the identity corresponding to the strongest authentication method where a choice is available.

3.7. Management

From an operational point of view, a network device in control of network advertisement and providing "realm hints" for guiding the network discovery and selection, should at least offer a management interface capable of providing status information for operators. Status information, such as counters of each selected network and used realm, and when RFC 4284 is used, the count of delivered "realm hints" might interest operators. Especially the information related to realms that fall into the "default free zone" or the "AAA fails to route" are of interest.

Larger deployments would benefit from a management interface that allow full remote configuration capabilities, for example, of "realm

hints" in case of RFC 4284-conforming network devices. While changes to "realm hints" and realm routing information are not expected to be frequent, centralized remote management tends to lower the frequency of misconfigured devices.

4. Conclusions

This document describes the network selection and discovery problem. In the opinion of the authors, the major findings are as follows:

- o There is a need for additional work on access network discovery, identifier selection, AAA routing, and payload routing.
- o Credential selection and AAA routing are aspects of the same problem, namely identity selection.
- o When considering selection among a large number of potential access networks and points of attachment, the issues described in the document become much harder to solve in an automated way, particularly if there are constraints on handoff latency.
- o The proliferation of network discovery technologies within IEEE 802, IETF, and 3rd Generation Partnership Project (3GPP) has the potential to become a significant problem going forward. Without a unified approach, multiple non-interoperable solutions may be deployed.
- o New link-layer designs should include efficient distribution of network and realm information as a design requirement.
- o It may not be possible to solve all aspects of the problem for legacy NAS devices on existing link layers. Therefore, a phased approach may be more realistic. For example, a partial solution could be made available for existing link layers, with a more complete solution requiring support for link layer extensions.

With respect to specific mechanisms for access network discovery and selection:

- o Studies such as [MACScale] and [Velayos], as well as the calculations described in Section 2.1, demonstrate that the IEEE 802.11 Beacon/Probe Response mechanism has substantial scaling issues in situations where a new Beacon is used for each "virtual AP". As a result, a single channel is, in practice, limited to less than twenty Beacon announcements with IEEE 802.11b.

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

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

As a result, there appears to be a need to enhance the scalability of IEEE 802.11 network advertisements.

- o Work is underway in IEEE 802.1, IEEE 802.21, and IEEE 802.11u [IEEE.802.11u] to provide enhanced discovery functionality. Similarly, IEEE 802.1af [IEEE.802.1af] has discussed the addition of network discovery functionality to IEEE 802.1X [IEEE.8021X-2004]. However, neither IEEE 802.1AB [IEEE.802.1ab] nor IEEE 802.1af is likely to support fragmentation of network advertisement frames so that the amount of data that can be transported will be limited.
- o While IEEE 802.11k [IEEE.802.11k] provides support for the Neighbor Report, this only provides for gathering of information on neighboring 802.11 APs, not points of attachment supporting other link layers. Solution to this problem would appear to require coordination across IEEE 802 as well as between standards bodies.
- o Given that EAP does not support fragmentation of EAP-Request/Identity packets, the volume of "realm hints" that can be fit with these packets is limited. In addition, within IEEE 802.11, EAP packets can only be exchanged within State 3 (associated and authenticated). As a result, use of EAP for realm discovery may result in significant delays. The extension of the realm advertisement mechanism defined in [RFC4284] to handle advertisement of realm capability information (such as QoS provisioning) is not recommended due to semantic and packet size limitations [GROETING]. As a result, we believe that extending the mechanism described in [RFC4284] for discovery of realm capabilities is inappropriate. Instead, we believe it is more appropriate for this functionality to be handled within the link layer so that the information can be available early in the handoff process.
- o Where link-layer approaches are not available, higher-layer approaches can be considered. A limitation of higher-layer solutions is that they can only optimize the movement of already connected hosts, but cannot address scenarios where network discovery is required for successful attachment.

Higher-layer alternatives worth considering include the SEAMOBY CARD protocol [RFC4066], which enables advertisement of network device capabilities over IP, and Device Discovery Protocol (DDP) [MARQUES], which provides functionality equivalent to IEEE 802.1AB using ASN.1 encoded advertisements sent to a link-local scope multicast address.

5. 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 to which the user tries to attach. Unfortunately, current EAP methods do not always make the verification of such parameters possible. EAP methods, such as Protected EAP (PEAP) [JOSEFSSON] and EAP-IKEv2 [IKEV2], may make this possible, however. There is even an attempt to provide a backward-compatible extension to older methods [ARKKO].

The security requirements for network selection depend on whether the selection is considered a mandate or a hint. In general, treating network advertisements as a hint is a more secure approach, since it reduces access client vulnerability to forged network advertisements. For example, "realm hints" may be ignored by an EAP peer if they are incompatible with the security policy corresponding to a selected access network.

Similarly, network access clients may refuse to connect to a point of attachment if the advertised security capabilities do not match those that have been pre-configured. For example, if an IEEE 802.11 access client has been pre-configured to require WPA2 enterprise support within an access network, it may refuse to connect to access points advertising support for WEP.

Where the use of methods that do not satisfy the security requirements of [RFC4017] is allowed, it may be possible for an attacker to trick a peer into using an insecure EAP method, leading to the compromise of long-term credentials. This can occur either where a network is pre-configured to allow use of an insecure EAP method, or where connection without pre-configuration is permitted using such methods.

For example, an attacker can spoof a network advertisement, possibly downgrading the advertised security capabilities. The rogue access point would then attempt to negotiate an insecure EAP method. Such

an attack can be prevented if the peer refuses to connect to access points not meeting its security requirements, which would include requiring use of EAP methods satisfying the [RFC4017] requirements.

Support for secure discovery could potentially protect against spoofing of network advertisements, enabling verifiable information to guide connection decisions. However, development of these mechanisms requires solving several difficult engineering and deployment problems.

Since discovery is a prerequisite for authentication, it is not possible to protect initial discovery using dynamic keys derived in the authentication process. On the other hand, integrity protection of network advertisements utilizing symmetric keys or digital signatures would require pre-configuration.

6. Informative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.
- [RFC3017] Riegel, M. and G. Zorn, "XML DTD for Roaming Access Phone Book", RFC 3017, December 2000.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [RFC4334] Housley, R. and T. Moore, "Certificate Extensions and Attributes Supporting Authentication in Point-to-Point Protocol (PPP) and Wireless Local Area Networks (WLAN)", RFC 4334, February 2006.
- [RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, December 2005.
- [RFC3280] Housley, R., Polk, W., Ford, W., and D. Solo, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3280, April 2002.

- [RFC4072] Eronen, P., Hiller, T., and G. Zorn, "Diameter Extensible Authentication Protocol (EAP) Application", RFC 4072, August 2005.
- [RFC3579] Aboba, B. and P. Calhoun, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)", RFC 3579, September 2003.
- [RFC2194] Aboba, B., Lu, J., Alsop, J., Ding, J., and W. Wang, "Review of Roaming Implementations", RFC 2194, September 1997.
- [RFC2607] Aboba, B. and J. Vollbrecht, "Proxy Chaining and Policy Implementation in Roaming", RFC 2607, June 1999.
- [RFC2608] Guttman, E., Perkins, C., Veizades, J., and M. Day, "Service Location Protocol, Version 2", RFC 2608, June 1999.
- [RFC3580] Congdon, P., Aboba, B., Smith, A., Zorn, G., and J. Roese, "IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines", RFC 3580, September 2003.
- [RFC4284] Adrangi, F., Lortz, V., Bari, F., and P. Eronen, "Identity Selection Hints for the Extensible Authentication Protocol (EAP)", RFC 4284, January 2006.
- [RFC4017] Stanley, D., Walker, J., and B. Aboba, "Extensible Authentication Protocol (EAP) Method Requirements for Wireless LANs", RFC 4017, March 2005.
- [RFC2486] Aboba, B. and M. Beadles, "The Network Access Identifier", RFC 2486, January 1999.
- [RFC4066] Liebsch, M., Singh, A., Chaskar, H., Funato, D., and E. Shim, "Candidate Access Router Discovery (CARD)", RFC 4066, July 2005.
- [IKEV2] Tschofenig, H., Kroeselberg, D., Pashalidis, A., Ohba, Y., and F. Bersani, "EAP-IKEv2 Method", Work in Progress, September 2007.
- [ARKKO] Arkko, J. and P. Eronen, "Authenticated Service Information for the Extensible Authentication Protocol (EAP)", Work in Progress, October 2005.

- [GROETING] Groeting, W., Berg, S., Tschofenig, H., and M. Ness, "Network Selection Implementation Results", Work in Progress, July 2004.
- [JOSEFSSON] Palekar, A., Simon, D., Salowey, J., Zhou, H., Zorn, G., and S. Josefsson, "Protected EAP Protocol (PEAP) Version 2", Work in Progress, October 2004.
- [MARQUES] Enns, R., Marques, P., and D. Morrell, "Device Discovery Protocol (DDP)", Work in Progress, May 2003.
- [OHBA] Taniuchi, K., Ohba, Y., and D. Subir, "IEEE 802.21 Basic Schema", Work in Progress, October 2007.
- [IEEE.802.11-2003] IEEE, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Standard 802.11, 2003.
- [Fixingapsell] Judd, G. and P. Steenkiste, "Fixing 802.11 Access Point Selection", Sigcomm Poster Session 2002.
- [IEEE.802.11k] IEEE, "Draft Ammendment to Standard for Telecommunications and Information Exchange Between Systems - LAN/MAN Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Radio Resource Management", IEEE 802.11k, D7.0, January 2007.
- [IEEE.802.1ab] IEEE, "Draft Standard for Local and Metropolitan Area Networks - Station and Media Access Control Connectivity Discovery", IEEE 802.1AB, D1.0, April 2007.
- [IEEE.802.1af] IEEE, "Draft Standard for Local and Metropolitan Area Networks - Port-Based Network Access Control - Amendment 1: Authenticated Key Agreement for Media Access Control (MAC) Security", IEEE 802.1af, D1.2, January 2007.

- [IEEE.802.11v]
IEEE, "Draft Amemdment to Standard for Information Technology - Telecommunications and Information Exchange Between Systems - LAN/MAN Specific Requirements - Part 11: Wireless Medium Access Control (MAC) and physical layer (PHY) specifications: Wireless Network Management", IEEE 802.11v, D0.09, March 2007.
- [Eronen04]
Eronen, P. and J. Arkko, "Role of authorization in wireless network security", Extended abstract presented in the DIMACS workshop, November 2004.
- [IEEE.11-04-0624]
Berg, S., "Information to Support Network Selection", IEEE Contribution 11-04-0624 2004.
- [Priest04]
Priest, J., "The State of Wireless London", July 2004.
- [MACScale]
Heusse, M., "Performance Anomaly of 802.11b", LSR-IMAG Laboratory, Grenoble, France, IEEE Infocom 2003.
- [Velayos]
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.
- [IEEE.802.11u]
IEEE, "Draft Amendment to STANDARD FOR Information Technology - LAN/MAN Specific Requirements - Part 11: Interworking with External Networks; Draft Amendment to Standard; IEEE P802.11u/D0.04", IEEE 802.11u, D0.04, April 2007.
- [IEEE-11-03-154r1]
Aboba, B., "Virtual Access Points", IEEE Contribution 11-03-154r1, May 2003.
- [IEEE-11-03-0827]
Hepworth, E., "Co-existence of Different Authentication Models", IEEE Contribution 11-03-0827 2003.

- [11-05-0822-03-000u-tgu-requirements]
Moreton, M., "TGu Requirements", IEEE Contribution 11-05-0822-03-000u-tgu-requirements, August 2005.
- [3GPPSA2WLANTS]
3GPP, "3GPP System to Wireless Local Area Network (WLAN) interworking; System Description; Release 6; Stage 2", 3GPP Technical Specification 23.234, September 2005.
- [3GPP-SA3-030736]
Ericsson, "Security of EAP and SSID based network advertisements", 3GPP Contribution S3-030736, November 2003.
- [3GPP.23.122]
3GPP, "Non-Access-Stratum (NAS) functions related to Mobile Station (MS) in idle mode", 3GPP TS 23.122 6.5.0, October 2005.
- [WWRF-ANS]
Eijk, R., Brok, J., Bemmels, 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.
- [WLAN3G]
Ahmavaara, K., Haverinen, H., and R. Pichna, "Interworking Architecture between WLAN and 3G Systems", IEEE Communications Magazine, November 2003.
- [INTELe2e]
Intel, "Wireless LAN (WLAN) End to End Guidelines for Enterprises and Public Hotspot Service Providers", November 2003.
- [Eronen03]
Eronen, P., "Network Selection Issues", presentation to EAP WG at IETF 58, November 2003.
- [3GPPSA3WLANTS]
3GPP, "3GPP Technical Specification Group Service and System Aspects; 3G Security; Wireless Local Area Network (WLAN) interworking security (Release 6); Stage 2", 3GPP Technical Specification 33.234 v 6.6.0, October 2005.

[3GPPCT1WLANTS]

3GPP, "3GPP System to Wireless Local Area Network (WLAN) interworking; User Equipment (UE) to network protocols; Stage 3 (Release 6)", 3GPP Technical Specification 24.234 v 6.4.0, October 2005.

[IEEE.802.21]

IEEE, "Draft IEEE Standard for Local and Metropolitan Area Networks: Media Independent Handover Services", IEEE 802.21, D05.00, April 2007.

[3GPPCT4WLANTS]

3GPP, "3GPP system to Wireless Local Area Network (WLAN) interworking; Stage 3 (Release 6)", 3GPP Technical Specification 29.234 v 6.4.0, October 2005.

[IEEE.8021X-2004]

IEEE, "Local and Metropolitan Area Networks: Port-Based Network Access Control", IEEE Standard 802.1X, July 2004.

Appendix A. Existing Work

A.1. IETF

Several IETF WGs have dealt with aspects of the network selection problem, including the AAA, EAP, PPP, RADIUS, ROAMOPS, and RADEXT WGs.

ROAMOPS WG developed the NAI, originally defined in [RFC2486], and subsequently updated in [RFC4282]. Initial roaming implementations are described in [RFC2194], and the use of proxies in roaming is addressed in [RFC2607]. The SEAMOBY WG developed CARD [RFC4066], which assists in discovery of suitable base stations. PKIX WG produced [RFC3280], which addresses issues of certificate selection. The AAA WG developed more sophisticated access routing, authentication, and service discovery mechanisms within Diameter [RFC3588].

Adrangi et al. [RFC4284] defines the use of the EAP-Request/Identity to provide "realm hints" useful for identity selection. The NAI syntax described in [RFC4282] enables the construction of source routes. Together, these mechanisms enable the user to determine whether it possesses an identity and corresponding credential suitable for use with an EAP-capable NAS. This is particularly useful in situations where the lower layer provides limited information (such as in wired IEEE 802 networks where IEEE 802.1X currently does not provide for advertisement of networks and their capabilities).

However, advertisement mechanisms based on the use of the EAP-Request/Identity have scalability problems. As noted in [RFC3748] Section 3.1, the minimum EAP Maximum Transmission Unit (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 [RFC1035] enables Fully Qualified Domain Names (FQDN) to be up to 255 octets in length, this may not enable the announcement of many realms. The use of network identifiers other than domain names is also possible.

As noted in [Eronen03], the use of the EAP-Request/Identity for realm 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 realms are supported. Since IEEE 802.11-2003 does not support use of Class 1 data frames in State 1 (unauthenticated, unassociated) within an Extended Service Set (ESS), this implies either that the APs must support 802.1X pre-authentication (optional in IEEE 802.11i-2004), or that the station must associate with each

AP prior to sending an EAPOL-Start to initiate EAP (here, EAPOL refers to EAP over LAN). This will dramatically increase handoff latency.

Thus, rather than thinking of [RFC4284] as an effective network discovery mechanism, it is perhaps better to consider the use of "realm hints" as an error recovery technique to be used to inform the EAP peer that AAA routing has failed, and perhaps to enable selection of an alternate identity that can enable successful authentication. Where "realm hints" are only provided in event of a problem, rather than as a staple network discovery technique, it is probably best to enable "realm hints" to be sent by core AAA proxies in the "default free" zone. This way, it will not be necessary for NASes to send "realm hints", which would require them to maintain a complete and up-to-date realm routing table, something that cannot be easily accomplished given the existing state of AAA routing technology.

If realm routing tables are manually configured on the NAS, then changes in the "default free" realm routing table will not automatically be reflected in the realm list advertised by the NAS. As a result, a realm advertised by the NAS might not, in fact, be reachable, or the NAS might neglect to advertise one or more realms that were reachable. This could result in multiple EAP-Identity exchanges, with the initial set of "realm hints" supplied by the NAS subsequently updated by "realm hints" provided by a core AAA proxy. In general, originating "realm hints" on core AAA proxies appears to be a more sound approach, since it provides for "fate sharing" -- generation of "realm hints" by the same entity (the core AAA proxy) that will eventually need to route the request based on the hints. This approach is also preferred from a management perspective, since only core AAA proxies would need to be updated; no updates would be required to NAS devices.

A.2. IEEE 802

There has been work in several IEEE 802 working groups relating to network discovery:

- o [IEEE.802.11-2003] defines the Beacon and Probe Response mechanisms within IEEE 802.11. Unfortunately, Beacons may be sent only at a rate within the base rate set, which typically consists of the lowest supported rate, or perhaps the next lowest rate. Studies such as [MACScale] have identified MAC layer performance problems, and [Velayos] has identified scaling issues from a lowering of the Beacon interval.
- o [IEEE-11-03-0827] 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 compare these scalability issues to those of alternative solutions, however.

- o [IEEE-11-03-154r1] 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 a 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 paper 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.
- o IEEE 802.11u is working on realm discovery and network selection [11-05-0822-03-000u-tgu-requirements] [IEEE.802.11u]. This includes a mechanism for enabling a station to determine the identities it can use to authenticate to an access network, prior to associating with that network. As noted earlier, solving this problem requires the AP to maintain an up-to-date, "default free" realm routing table, which is not feasible without dynamic routing support within the AAA infrastructure. Similarly, a priori discovery of features supported within home realms (such as enrollment) is also difficult to implement in a scalable way, absent support for dynamic routing. Determination of network capabilities (such as QoS support) is considerably simpler, since these depend solely on the hardware and software contained within the AP. However, 802.11u is working on Generic Advertisement Service (GAS) mechanism, which can be used to carry 802.21 Information Service (IS) messages and, in that way, allow a more sophisticated way of delivering information from the network side.
- o IEEE 802.21 [IEEE.802.21] is developing standards to enable handover between heterogeneous link layers, including both IEEE 802 and non-IEEE 802 networks. To enable this, a general mechanism for capability advertisement is being developed, which could conceivably benefit aspects of the network selection problem, such as realm discovery. For example, IEEE 802.21 is developing Information Elements (IEs) that may assist with network selection, including information relevant to both layer 2 and layer 3. Query mechanisms (including both XML and TLV support) are also under development. IEEE 802.21 also defines a Resource Description Framework (RDF) schema to allow use of a query

language (i.e., SPARQL). The schema is a normative part of IEEE 802.21 and also defined in [OHBA].

A.3. 3GPP

The 3GPP stage 2 technical specification [3GPPSA2WLANTS] covers the architecture of 3GPP Interworking WLAN (I-WLAN) with 2G and 3G networks. This specification also discusses realm discovery and network selection issues. The I-WLAN realm discovery procedure borrows ideas from the cellular Public Land-based Mobile Network (PLMN) selection principles, known as "PLMN Selection".

In 3GPP PLMN selection [3GPP.23.122], the mobile node monitors surrounding cells and prioritizes them based on signal strength before selecting a new potential target cell. Each cell broadcasts its PLMN. A mobile node may automatically select cells that belong to its Home PLMN, Registered PLMN, or an allowed set of Visited PLMNs. The PLMN lists are prioritized and stored in the Subscriber Identity Module (SIM). In the case of manual PLMN selection, the mobile node lists the PLMNs it learns about from surrounding cells and enables the user to choose the desired PLMN. After the PLMN has been selected, cell prioritization takes place in order to select the appropriate target cell.

[WLAN3G] discusses the new realm (PLMN) selection requirements introduced by I-WLAN roaming, which support automatic PLMN selection, not just manual selection. Multiple network levels may be present, and the hotspot owner may have a contract with a provider who, in turn, has a contract with a 3G network, which may have a roaming agreement with other networks.

The I-WLAN specification requires that network discovery be performed as specified in the relevant WLAN link layer standards. In addition to network discovery, it is necessary to select intermediary realms to enable construction of source routes. In 3GPP, the intermediary networks are PLMNs, and it is assumed that an access network may have a roaming agreement with more than one PLMN. The PLMN may be a Home PLMN (HPLMN) or a Visited PLMN (VPLMN), where roaming is supported. GSM/UMTS roaming principles are employed for routing AAA requests from the VPLMN to the Home Public Land-based Mobile Network (HPLMN) using either RADIUS or Diameter. The procedure for selecting the intermediary network has been specified in the stage 3 technical specifications [3GPPCT1WLANTS] and [3GPPCT4WLANTS].

In order to select the PLMN, the following procedure is required:

- o The user may choose the desired HPLMN or VPLMN manually or let the WLAN User Equipment (WLAN UE) choose the PLMN automatically, based on user and operator defined preferences.
- o AAA messages are routed based on the decorated or undecorated NAI.
- o EAP is utilized as defined in [RFC3748] and [RFC3579].
- o PLMN advertisement and selection is based on [RFC4284], which defines only realm advertisement. The document refers to the potential need for extensibility, though EAP MTU restrictions make this difficult.

The I-WLAN specification states that "realm hints" are only provided when an unreachable realm is encountered. Where VPLMN control is required, this is handled via NAI decoration. The station may manually trigger PLMN advertisement by including an unknown realm (known as the Alternative NAI) within the EAP-Response/Identity. A realm guaranteed not to be reachable within 3GPP networks is utilized for this purpose.

The I-WLAN security requirements are described in the 3GPP stage 3 technical specification [3GPPSA3WLANTS]. The security requirements for PLMN selection are discussed in 3GPP contribution [3GPP-SA3-030736], which concludes that both SSID and EAP-based mechanisms have similar security weaknesses. As a result, it recommends that PLMN advertisements should be considered as hints.

A.4. Other

[INTELe2e] discusses the need for realm selection where an access network may have more than one roaming relationship path to a home realm. It also describes solutions to the realm selection problem based on EAP, SSID and Protected EAP (PEAP) based mechanisms.

Eijk et al. [WRRF-ANS] discusses the realm and network selection problem. The authors concentrate primarily on discovery of access networks meeting a set of criteria, noting that information on the realm capabilities and reachability inherently resides in home AAA servers, and therefore it is not readily available in a central location, and may not be easily obtained by NAS devices.

Appendix B. Acknowledgements

The authors of this document would like to especially acknowledge the contributions of Farid Adrangi, Michael Richardson, Pasi Eronen, Mark Watson, Mark Grayson, Johan Rune, and Tomas Goldbeck-Lowe.

Input for the early versions of this document has been gathered from many sources, including the above persons as well as 3GPP and IEEE developments. We would also like to thank Alper Yegin, Victor Lortz, Stephen Hayes, and David Johnston for comments.

Jouni Korhonen would like to thank the Academy of Finland for providing funding to work on this document.

Authors' Addresses

Jari Arkko
Ericsson
Jorvas 02420
Finland

EMail: jari.arkko@ericsson.com

Bernard Aboba
Microsoft
One Microsoft Way
Redmond, WA 98052
USA

EMail: bernarda@microsoft.com

Jouni Korhonen
TeliaSonera
Teollisuuskatu 13
Sonera FIN-00051
Finland

EMail: jouni.korhonen@teliasonera.com

Farooq Bari
AT&T
7277 164th Avenue N.E.
Redmond WA 98052
USA

EMail: farooq.bari@att.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

