

The Internet Protocol Journal

June 2007

Volume 10, Number 2

A Quarterly Technical Publication for
Internet and Intranet Professionals

FROM THE EDITOR

In This Issue

From the Editor	1
AAA—Part Two	2
IPv6 Network Mobility	16
More ROAP	28
Time to Replace SMTP?	34
Fragments	39

Part One of a two-part article on *Authentication, Authorization, and Accounting* (AAA) was published in our previous issue. This time Sean Convery presents Part Two—subtitled “Protocols, Applications, and the Future of AAA.”

Interest in *IP Version 6* (IPv6) is growing in many parts of the Internet technical community; see, for example, the announcement from ARIN on page 39 of this issue. Transition to IPv6 is likely to be one of the greatest technical challenges in the history of the Internet. Several groups are developing parts of the overall solution by creating IPv6-capable versions of protocols such as the *Dynamic Host Configuration Protocol* (DHCP) or including support for IPv6 in the *Domain Name System* (DNS). Although not yet widely deployed, *IP Network Mobility* is expected to play an important part in the Internet of the future. For this reason the IETF is working on IP mobility with an eye toward IPv6. Our second article looks at the *Network Mobility (NEMO) Basic Support Protocol*, which is being developed by the NEMO working group in the IETF.

Depletion of IPv4 address space is not the only concern for network operators and developers these days. Questions about the long-term viability of today’s routing protocols and the associated addressing systems center around a basic concern about how we can scale our networks to a size orders of magnitude larger than what we have today. A recently formed *Routing and Addressing Problem Directorate* (ROAP) is tasked to examine these problems in detail. Several ROAP-related sessions took place during the most recent IETF meeting, and Geoff Huston reports on these sessions and gives his analysis and commentary. Incidentally, Geoff was not present in person at this IETF meeting, but the facilities to follow an IETF meeting remotely are now of such a quality that he was able to participate from the other side of the world.

Protocol replacement or enhancement is also the theme in our final article. Dave Crocker asks the question “Is it time to replace SMTP?” Since this is an opinion piece, we invite your feedback or rebuttals.

New on our Website is a linked article index. Visit cisco.com/ipj and click on “Index Files” to explore this feature.

—Ole J. Jacobsen, Editor and Publisher
ole@cisco.com

You can download IPJ
back issues and find
subscription information at:
www.cisco.com/ipj

Network Authentication, Authorization, and Accounting Part Two: Protocols, Applications, and the Future of AAA

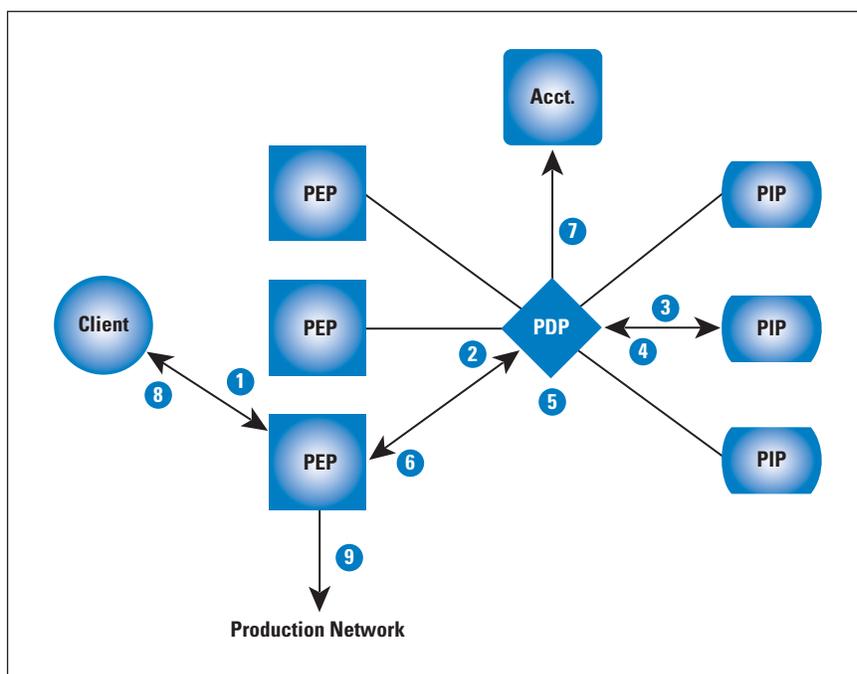
by Sean Convery, Identity Engines

Network Authentication, Authorization, and Accounting has been used since before the days of the Internet as we know it today. Authentication asks the question, “Who or what are you?” Authorization asks, “What are you allowed to do?” And finally, accounting wants to know, “What did you do?” These fundamental security building blocks are being used in expanded ways today. The first part of this two-part series focused on the overall concepts of AAA, the elements involved in AAA communications, and high-level approaches to achieving specific AAA goals. It was published in IPJ Volume 10, No. 1^[0]. This second part of the series discusses the protocols involved, specific applications of AAA, and considerations for the future of AAA.

AAA Protocols

Although AAA is often thought of as the exclusive province of the *Remote Authentication Dial-In User Service* (RADIUS) protocol, in reality a range of protocols is involved at various stages of the AAA conversation. This section introduces these AAA protocols, organized according to the parties involved in the communication. We divide AAA communications into the following categories: Client to *Policy Enforcement Point* (PEP), PEP to *Policy Decision Point* (PDP), Client to PDP, and PDP to *Policy Information Point* (PIP). For easy reference, the AAA flow diagram from Part One of this article is reproduced here. Please refer to Part One^[0] for the explanatory text associated with the diagram.

Figure 1: A Client Connects to a AAA-Protected Network (from Part One)



Client to PEP

AAA communications between the client and the PEP can travel at Layer 2 of the OSI model, or they can run at higher layers, relying on lower layers as essentially dumb transport. The most common protocols for client-to-PEP communication are the *Point-to-Point Protocol* (PPP)^[1], *PPP over Ethernet* (PPPoE)^[2], IEEE 802.1X^[3], *IP Security* (IPsec), *Secure Sockets Layer* (SSL) VPN, and *Hypertext Transfer Protocol* (HTTP), each of which is discussed in this article.

PPP, the standard protocol for communicating across point-to-point links, includes an optional authentication step—the point at which the AAA element is introduced. During this authentication phase, protocols such as the *Challenge Handshake Authentication Protocol* (CHAP) can be used to identify the client to the PEP. (These protocols were discussed in the credential section of Part One of this article.) PPP is extensively used in dialup access but is otherwise not found in modern AAA. PPPoE, an adaptation of PPP to run over Ethernet, is used by many service providers rolling out broadband services.

PPPoE allows the broadband endpoint to authenticate itself to the service provider’s network when making the initial connection. Because many broadband networks use shared Ethernet mediums, PPPoE allows *Internet Service Providers* (ISPs) to maintain the per-user accounting they were familiar with from dialup. The 802.1X protocol is an IEEE standard specifying a way to provide network access control at the port level for wired and wireless networks. The 802.1X standard specifies a way for the client to communicate with the PDP using the *Extensible Authentication Protocol* (EAP)^[4], which is discussed in more detail later in this section. The 802.1X standard requires that the endpoint support 802.1X through a “supplicant” or client sign-on application. This application authenticates the client to the network through the PEP. (See the EAP section later in this article for an explanation showing how EAP and 802.1X can work together.)

For wireless networks, 802.1X has become the standard way of authenticating clients because it supports communicating unique key material to the client to secure its use of the wireless infrastructure. In wired Ethernet networks, 802.1X is rising in popularity as a way to authenticate clients as well. These applications are more fully described in the “AAA Applications” section, later in this article.

At a more generic level, the IPsec protocol has established a standard for securing IP communications, and this approach has become another common method of communicating from a client to a PEP (referred to as a *VPN Gateway* from an IPsec perspective). The initial authentication for IPsec communications uses the *Internet Key Exchange* (IKE) protocol. Version 1^[5] of the IKE protocol had no built-in method for authenticating users with credentials such as passwords, so an extension to IKE called *XAUTH*^[6] was proposed.

XAUTH never became an official standard (though it certainly was a de facto one) because the IETF IPsec working group created a second version of IKE^[7] that used EAP as a transport for credentials such as passwords. Finally, in the areas of HTTP and VPN communications, the SSL and *Transport Layer Security* (TLS)^[28] standards are two closely related protocols for securing, among other things, Web communications. SSL/TLS VPNs use these protocols to create a secure session from the client to the PEP (VPN Gateway). Client authentication with SSL and TLS can be done with client-side certificates, but more commonly they use passwords or *One-Time Passwords* (OTPs).

PEP to PDP

The three main protocols for communicating between a PEP and a PDP are TACACS+^[9], RADIUS, and *Diameter*^[10]. First, consider TACACS+: Developed by Cisco, TACACS+ is a proprietary protocol that is used primarily in communicating administrator authorizations for network devices. TACACS+ uses TCP port 49 and features payload encryption for the entire TACACS+ message. Though developed by Cisco, TACACS+ is supported by other companies as well, including Juniper.

Although TACACS+ excels at command-level authorizations and accounting for administrator control, another protocol has become far more common for client AAA: RADIUS. Thanks to nearly ubiquitous support for this protocol in network hardware, RADIUS is the primary protocol for communication between a PEP and a PDP in most environments. RADIUS uses the *User Datagram Protocol* (UDP) port 1812 for authentication and authorization and UDP port 1813 for accounting^[8] (early deployments used ports 1645 and 1646, which are still used sometimes today). RADIUS supports numerous different attributes for communicating information back and forth from the PEP to the PDP, such as client MAC address, username, filter information for enforcement, and so on. It also supports an extensible framework for *Vendor-Specific Attributes* (VSAs), which allow extensions of the functions of RADIUS to support whatever elements a given PEP might need to best serve its role on the network. For example, a PEP manufacturer might support VSAs that allow the assignment of a user to a particular enforcement profile. RADIUS in its default implementation encrypts only the Password field of RADIUS messages, making the RADIUS protocol more prone to leaking information that could be used by an adversary. Both RADIUS and TACACS+ are secured by only a shared secret that is configured on both the PEP and the PDP.

Finally, consider the Diameter protocol. Diameter (the name is a play on words from RADIUS) is the next-generation, de jure standard for AAA. It supports stronger security through either IPsec or TLS and greater extensibility than RADIUS. It uses port 3868 for either TCP or the *Stream Control Transmission Protocol* (SCTP)^[11]. The strongest use of Diameter to date is in the carrier space, where it provides AAA for call processing and *third-generation* (3G) mobile networks.

However, the corporate market has been fairly reluctant to embrace Diameter, and that reluctance has translated into a lack of support for Diameter in corporate network infrastructure equipment.

At this point in the discussion, it makes sense to compare RADIUS and Diameter. Although Diameter is an obvious alternative, RADIUS continues to be used in both new and existing deployments, so the IETF has a working group specifically formed to extend RADIUS in the future. The relationship between RADIUS and Diameter is a little like the relationship between IPv4 and IPv6. IPv6 had IPsec as a standard feature, IPv4 integrated IPsec as well, and today, by a large margin, most IPsec deployments are on IPv4 networks. The situation is similar with AAA. RADIUS certainly had limitations, but since Diameter entered the picture, RADIUS has been extended to address some of those shortcomings, particularly with both protocols using EAP as a transport. The result is that RADIUS today does what most people want. Therefore, given the significant added complexity of Diameter, many organizations have elected not to migrate to Diameter. Both RADIUS and Diameter will be around for many years to come.

Client to PDP

Although most of the protocols in this article handle communication from one component to the next component in the AAA chain (that is, client to PEP, PEP to PDP, etc.), there is one protocol that deals with communication from the client to the PDP directly: the *Extensible Authentication Protocol* (EAP). As mentioned earlier, EAP is a flexible mechanism for communicating almost any kind of credential over almost any lower-layer transport. Each technique for authenticating a client is referred to as an *EAP Method*. Originally conceived as an extension to PPP, EAP can now use many transports, including IKEv2 and 802.1X. Cisco's proprietary *Network Admission Control* (NAC) solution offers a deployment option that puts EAP inside UDP. When using 802.1X, for example, EAP uses LAN transport, referred to as *EAPoL* (EAP over LAN). This transport is only for the connection between the client and the PEP though. From the PEP to the PDP, EAP rides inside RADIUS^[12, 13]. The actual conversation, however, takes place between the client and the PDP, with the PEP acting as a relay.

The major benefit of this approach is that the PEP does not need to understand the specifics of the EAP method selected—only the client and the PDP do. The EAP specification in the IETF specifies several different EAP methods, including *EAP Message Digest Algorithm 5* (EAP-MD5, very similar in security to CHAP), *EAP-OTP* (which supports an IETF-defined OTP solution^[14]), and *EAP Generic Token Card* (EAP-GTC). Of the methods explicitly called out in the EAP standard, EAP-GTC is the only one in much use today in production networks. EAP-GTC allows the use of OTP token cards within an EAP context.

Beyond the methods defined in the EAP standard, EAP by its nature can be extended to support additional methods. EAP *Subscriber Identity Module* (EAP-SIM)^[15] specifies a method for authentication using SIM elements in the *Global System for Mobile Communications* (GSM). EAP-SIM was developed by the *Third Generation Partnership Project* (3GPP) as a solution for these second-generation (GSM) mobile networks. EAP-AKA^[16] is the 3GPP's EAP authentication technique for third-generation (*Universal Mobile Telecommunications Service* [UMTS] or *Code Division Multiple Access 2000* [CDMA2000]) mobile networks. Both EAP-SIM and EAP-AKA support authenticating a mobile phone to a Wi-Fi network without using passwords. The problem is that without some sort of user identity federation solution in place, SIM-based authentication can work only with the mobile provider's network that supplied the SIM card. EAP-TLS^[17] specifies a technique for mutual certificate authentication. Although it is widely supported, EAP-TLS is not commonly deployed because of its requirement for client-side certificates.

Though none of the following EAP methods are standards, they—somewhat confusingly—represent the vast majority of EAP deployments. Each of them is referred to as a *Tunneled EAP Method* because it establishes one outer EAP method as a base secure channel and then runs another method (one that may be less secure) over that secure channel. *Protected EAP* (PEAP)^[18], well supported in Microsoft's Windows operating system, has become a de facto standard for EAP methods. Most clients and PDPs support PEAP today. PEAP works by establishing a TLS session authenticated by the server certificate, and then an inner authentication method rides inside that TLS session. The inner method is almost always *Microsoft CHAP Version 2* (MS-CHAPv2), but other methods can be used as well. Another popular tunneled protocol is *EAP Tunneled TLS* (EAP-TTLS)^[19]. This protocol is similar to PEAP except it supports a more arbitrary exchange of information inside the TLS tunnel. For example, one of the primary uses for EAP-TTLS is using the *Password Authentication Protocol* (PAP) as the inner authentication method, allowing an EAP-TTLS-capable PDP to authenticate clients against older password stores (such as those that support only PAP authentication).

Finally, in settings that use primarily Cisco equipment, a common tunneled protocol is *EAP Flexible Authentication via Secure Tunneling* (EAP-FAST)^[20]. This protocol uses TLS to authenticate the PDP, and then a shared key is distributed to allow faster subsequent authentication. An inner EAP method such as MS-CHAPv2 can then be used to authenticate the client to the server. EAP-FAST is used extensively in Cisco products for wireless deployments.

PDP to PIP

The final set of AAA protocols we consider are the ones that govern the communication between the PDP and the PIP. The primary protocol of interest is the *Lightweight Directory Access Protocol* (LDAP)^[21]. From a AAA context, LDAP allows a PDP to query a PIP (typically an X.500 directory^[22]) for information about a client. This information is exposed through a series of group and attribute identifiers, which can include information about a client's home location, organizational role, job title (if referring to a user), and so on. LDAP includes several different authentication options^[23]. This client information learned from the PIP enables the PDP to better make its policy decision. Also useful in the PDP-PIP communications context is the RADIUS protocol. Some large organizations or inter-organization federations use a hierarchy of RADIUS-speaking PDPs where one RADIUS PDP can act as a PIP for another RADIUS PDP further down the AAA hierarchy.

Finally, Microsoft *Active Directory* (AD) uses the LDAP protocol when acting as a PDP but also has its own extension, called *Netlogon*, for validating Microsoft credentials such as MS-CHAPv2. This means that integrating a PDP with Microsoft AD generally involves using LDAP to find information about the client and using Netlogon to validate the client's credential. Other options for PDP-to-PIP interaction—though less often used—include *Structured Query Language* (SQL) databases, *Network Information Service* (NIS), and Kerberos.

AAA Applications

This section surveys the different applications of AAA technology throughout networking. It is divided into three sections covering consumer, enterprise, and carrier applications, with a final section covering emerging applications of AAA technology.

Consumer-Managed Applications

Most consumer network deployments do not perform any advanced AAA beyond a shared key for authentication to a wireless network. In this example, the client is the consumer's host and the wireless access point acts as PEP, PDP, and PIP by validating that any client connecting to the access point presents the correct shared key.

Enterprise-Managed Applications

AAA has numerous enterprise applications, including remote access, wireless security, *Voice over IP* (VoIP), guest access, *Role-Based Access Control* (RBAC), and endpoint posture validation (also known as NAC). This section discusses each of these applications. Remote-access security is the original enterprise AAA application. In the remote-access scenario, remote users connect over a dialup connection or a VPN and authenticate themselves (and optionally their hosts) to the organization's network.

The client's credential is almost always a password, expressed in one of the forms discussed in the credential section of Part One of this article. The main purpose of AAA in the remote-access case is to validate that the client is a valid user of the organization's network.

Wireless security is similar in some respects to remote-access security. The goals of AAA in wireless security are twofold: first it must validate that the wireless client is an authorized user, and second, it must provide the client with a session key for cryptographic protection of the client's traffic. Given these goals, 802.1X using EAP are the ideal protocols to use because they support both client authentication and dynamic keying. Older wireless security approaches relied on an open wireless network and a VPN Gateway separating that network from the rest of the organization's network. In that example, the wireless-security approach mimics the remote-access application just discussed. Other types of networking require different applications of AAA. For example, VoIP deployments have authentication requirements as well. The *Session Initiation Protocol* (SIP)^[24] is used extensively for, well, session initiation in VoIP networks (for example, authenticating the calling parties prior to initiating a new call). Authentication can be handled natively within SIP using HTTP digest authentication, or the same request can be sent to a PDP using RADIUS. AAA for VoIP allows handsets to authenticate themselves to the network and gain access to call-processing services.

Another, very popular application of AAA is guest-access management for networks. This application has grown quickly with the recent growth of wireless networks. Guest access is a method by which guests can be granted temporary access to a network with a full audit trail^[27]. Guest access generally involves creating a distinct PIP, which houses short-term user accounts, and a technique for creating and, after a configurable period of time, automatically deactivating those user accounts. The PIP is often co-resident with the PDP and allows this temporary access without having to provision these users into the organization's more permanent directory. The guest can communicate with the PEP using any of the client-PEP protocols discussed earlier, though HTTP is the most common. The PEP is told by the PDP that the client (because it is a guest) should have restricted access—typically access only to the Internet at large and not any communication with an organization's internal network.

Also growing in popularity as a AAA application is RBAC, an application of AAA that allows customization of the network session based on the role of the client. In fact, guest access is a simple form of RBAC whereby two classes of clients are created: guest and permanent. However, RBAC can be extended to include more levels of delineation, including guest, contractor, and specific classes of permanent users such as sales, human resources, and engineering.

This classification can be done with all forms of AAA-enabled network infrastructure, including wired, wireless, and remote access. Current scalability limitations of VLAN technology and *Access Control Lists* (ACLs) make creating large quantities of roles difficult, but a significant business benefit in audit and regulatory compliance can be realized with usually fewer than five roles.

To implement RBAC, most organizations choose a mix of 802.1X and HTTP authentication for wired and wireless access, combined with VPN technology for remote access. This approach is the most common one to RBAC, though others are used.

Finally, another important AAA application is *Endpoint Posture Validation*, also referred to as *Network Access Control* (NAC). Unfortunately NAC is an inappropriate name because of its almost complete overlap with the more general AAA term—leading to a fair amount of confusion in the market. Endpoint posture validation refers to many different parameters in the industry as it is an emerging technology. These parameters range from very narrow device-centric posture checking to a more identity-centric approach for secure mobile computing. Because this entire article is concerned with the latter, we will consider NAC in its narrow context of endpoint posture checking. With this label, NAC simply acts as another PIP for the PDP to use.

This time, though, instead of checking the client's credential, NAC checks the client's software configuration. This checking generally focuses on security-sensitive configuration details of the endpoint security software and the operating system itself, such as the revision, configuration, and current operating status. This client configuration data is gathered by a host agent on the client and then sent to the PDP or PIP for evaluation. The host agent is either permanent on the client or downloaded dynamically to acquire the information. Some NAC applications rely exclusively on external scanning of the client, although this scanning generally yields far less granular information than an agent would.

The challenge with NAC today is deploying a system built on standards. The IETF and the *Trusted Computing Group* (TCG) are both pursuing standards in this space. Meanwhile Cisco, Microsoft, and a host of smaller companies have offerings not currently based on any standard. Recent announcements from the TCG and Microsoft are changing this. The TCG recently standardized the as-implemented NAC protocol used by Microsoft's NAC approach. Though there is much more work to do, this should allow the beginnings of standards-based interoperability in NAC solutions since a core protocol in Microsoft's NAC is now a standard from the TCG. There is a great base in standards at a low enough layer in all the NAC approaches though, as the emerging standards use the protocols discussed in this article including 802.1X, IPsec, RADIUS, and LDAP.

Carrier-Managed Applications

Some carrier-managed AAA applications are similar to those for the enterprise and others are different. The common distinctions for almost all carrier applications are their large scale and their emphasis on accounting. Carrier applications include dialup, DSL or cable PPPoE, mobile or 3G, wireless hotspot, and metro wireless. Dialup is similar to the remote-access application in the enterprise section, but on a massive scale. *Network Access Servers* (NASs) for a large ISP are geographically dispersed, as are the PDP and PIP systems that support them. Clients communicate with the PEP (NAS) with PPP using one of the password credential techniques discussed in Part One of this article, and the PEP communicates with the PDP using RADIUS or Diameter.

Now consider DSL or cable PPPoE. Though PPPoE-based broadband access seems to be on the decline, many ISPs are still using PPPoE for the enhanced audit trail it provides compared with an unauthenticated connection. In the realm of mobile telephone networks, service providers are increasingly providing data services in mobile phones, and these services require AAA for security and billing. Such data services come in several varieties on both the second- and third-generation mobile networks. Additionally, smartphones are increasingly supporting 802.11-based wireless access as well, creating a complex relationship between the smartphone, mobile voice network, mobile data network, 802.11 data network, and VoIP-based voice services. Previously discussed standards such as EAP-SIM and EAP-AKA are trying to bridge some of these worlds, but there is much work to be done. Ideally, any smartphone should take advantage of the network with the fastest and richest set of services, and callers trying to reach a smartphone user as well as the user himself, should be shielded from this discovery and association process. Business motivators and detractors within the carrier space may affect this convergence.

The next carrier-managed AAA application to discuss is the *wireless hotspot*. Hotspots work much like dialup providers in that regular users get a password-based credential that lets them authenticate to the hotspot. In this context, the 802.1X protocol is less commonly used because the required client software is not yet ubiquitous in the client install base. More common is Web-based authentication much like that used to access broadband in a hotel. A critical security consideration for a hotspot operator is the ability to ensure that a given client is not connected to two hotspots at the same time—a situation that would indicate an account was shared between two or more users. This stipulation places an increasing burden on the accounting aspect of AAA, as with any carrier-based AAA application.

Finally, the last AAA application we examine is the metropolitan wireless network, known as “metro wireless.” In metro wireless, an 802.11 network is deployed throughout a metropolitan area, and access is provided free of charge or for a fee. I live in Mountain View, California, which is home to Google’s headquarters, and is where Google has installed its free, citywide metro wireless network.

Although the service is free, AAA is still required: to sign on to the wireless network, you must authenticate to Google using an ID. This step, much like signing on to a wireless hotspot, allows Google to trace network use to an individual (if necessary) and switch to a fee-based model later on if desired. HTTP authentication is most common in metro wireless environments, and, because of the on/off nature of access, little sophistication in policy decision is required other than validating the client's credential.

Emerging Applications

Several interesting applications of network AAA are emerging. The first is in building just-in-time networks, such as when establishing an on-scene emergency operations center after a disaster. In this situation, emergency workers often need to communicate in a protected environment, and the press that covers the disaster needs network access to send in its reports. The AAA application required here is a cross between wireless security, guest access, and RBAC.

Another emerging application is what we call "granular RBAC." As opposed to RBAC, which associates users into coarse-grained classes of users, granular RBAC knows much more about the users and makes a more sophisticated access decision.

One example of the use of granular RBAC is for classroom control in higher education. Increasingly, classrooms are wireless-enabled as a convenience feature for faculty and students. However, during exam time it is often useful to disable this access to the students taking an exam. Without a granular understanding of which clients are connecting to the network, this setup is very difficult to achieve without physically disabling large portions of the wireless network during exam time. By using AAA, a school could put class schedules inside an LDAP store along with the rest of the students' information. Professors could also register exam times by time and location. AAA could then prevent students from getting on the network inside the classroom during their exam period, while still letting them connect to the network when inside their dorm room.

Finally, the last application we consider is what I call "punitive access restrictions." As networks become more and more an integral part of our lives, it is natural to want as fast a network connection as we can find, creating the situation where denying access to the network based on past behavior (network related or not) can be used as a punitive action. Today, your driver's license can be revoked based on your behavior while on the road. Punitive access restrictions on the network could mirror the same technique (for example, punishing people who propagate a virus by restricting their network access for a time) or could be used even if the infraction is not related to the network. Imagine a university that has trouble getting students to return overdue library books. Fines are one way to get the books back quickly, but if the student's parents are paying the bill, this consequence may not be as effective as the university desires.

However, imagine if the student's account record (in the PIP) had a directory attribute containing a count of the student's overdue library books. The network could then use RBAC or *Quality-of-Service* (QoS) techniques to provide degraded access to the student until the books were returned.

The Future of AAA

AAA as a concept has remained relatively unchanged since its inception. However, as this article has demonstrated, the techniques and applications of AAA continue to evolve. This section discusses some of the ways AAA may change more fundamentally in the future.

Security and Identity Convergence

Today the security and identity services provided by physical building access, network access, and application access are completely distinct. Security can be improved by communicating among these layers. Imagine a user executing a \$10 million purchase order in a financial application. The chance of fraud would be reduced if the application could know that the user was coming from an authorized client with an up-to-date antivirus configuration. The chance of fraud could be further reduced by checking that the same user had accessed the badge access system of the building that day, and that the point of badge-access entry was consistent with the location where the application request originated. Within computer security, the notion of *defense-in-depth* has been around for a long time and is considered a best practice. Security and identity convergence adds new layers to this defense, and can potentially make all the layers more intelligent in their interaction.

User-Centric AAA

In the Web application world, the notion of user-centric identity is gaining ground. Kim Cameron's "Laws of Identity"^[25] makes a compelling case that identity information housed in silos to be used by one organization is problematic. Several circles in the Web and e-commerce communities are beginning to look at identity differently. One change, consistent with the notion of user-centric identity, is that users should own their own identity information and should control how that information is used. The simplest example I can offer is shopping preferences at an online store. Most online stores make suggestions to you based on prior purchases. This data is owned by the online store, though, and not you, the consumer. If you wanted to take your purchasing profile from Amazon.com and transfer it to Barnes and Noble, it would not work. With user-centric identity, this kind of process is possible.

Another example is asserting a user's age. Depending on what you are trying to do on the Internet, you may need to validate that you are above a certain age. To do that, you are often asked to enter your date of birth, but that is more information than the site really needs.

If you could assert, with an identity you control but that is validated by a trusted party, that you are over the required age, it would not be necessary to disclose your date of birth (a process that is sometimes used as an authentication factor when you call places such as your credit card company).

The idea here is that you control your own information and limit what you need to share with others. This is very beneficial for privacy. One user-centric identity approach is included in Windows Vista through an application called “Card Space.” Other approaches include OpenID and the Higgins Project. All of these approaches are somewhat consumer-focused, but if they take hold, it seems natural that there will be pressure for similar identity approaches in the enterprise and carrier space.

Federation

One of the natural evolutions of AAA infrastructure is to start federating access between multiple organizations. Imagine if visiting professors at another university’s campus could access the network as guests using their password from their home location? Federation promises to make this possible, but the most challenging hurdles are political and logistical rather than technological. Protocols such as the *Security Assertion Markup Language* (SAML)^[26] combined with RADIUS and LDAP can overcome this hurdle. The challenge is how to set up the trust relationships between the organizations to make it work. Eduroam based on RADIUS is an early effort delivering federation in Europe today.

Summary

This article, with its companion piece, has explored all aspects of AAA. Part One described the overall approach of AAA, how it works, and the elements that provide authentication, authorization, and accounting. Part Two has explored all the protocols used in the communication between the various AAA elements, the applications of AAA, and some thoughts about the future of AAA. AAA is a giant topic, and each of these sections, protocol descriptions, and applications could be expanded into a paper all by itself. The information in this article, combined with the references provided, should be a good starting point for your own examination of the specific aspects of AAA that are of interest to you.

References

- [0] Convery, S., “Network Authentication, Authorization, and Accounting —Part One: Concepts, Elements, and Approaches,” *The Internet Protocol Journal*, Volume 10, No. 1, March 2007.
- [1] Simpson, W., “The Point-to-Point Protocol (PPP),” RFC 1661, July 1994.

- [2] Mamakos, L., “A Method for Transmitting PPP Over Ethernet (PPPoE),” RFC 2516, February 1999.
- [3] Jeffree et al., “Port-Based Network Access Control,” IEEE Std 802.1X-2004, November 2004.
- [4] Aboba et al., “Extensible Authentication Protocol,” RFC 3748, June 2004.
- [5] Harkins et al., “The Internet Key Exchange (IKE),” RFC 2409, November 1998.
- [6] Beaulieu et al., “Extended Authentication within IKE (XAUTH),” Internet Draft, Work in Progress, October 2001.
draft-beaulieu-ike-xauth-02.txt
- [7] Kaufman C., ed., “Internet Key Exchange (IKEv2) Protocol,” RFC 4306, December 2005.
- [8] Rigney C., “RADIUS Accounting,” RFC 2866, June 2000.
- [9] Carrel et al., “The TACACS+ Protocol Version 1.78,” Internet Draft, Work in Progress, January 1997.
draft-grant-tacacs-02.txt
- [10] Calhoun et al., “Diameter Base Protocol,” RFC 3588, September 2003.
- [11] Stewart et al., “Stream Control Transmission Protocol,” RFC 2960, October 2000.
- [12] Aboba et al., “RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP),” RFC 3579, September 2003.
- [13] Congdon et al., “IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines,” RFC 3580, September 2003.
- [14] Haller et al., “A One-Time Password System,” RFC 2289, February 1998.
- [15] Haverinen et al., “Extensible Authentication Protocol Method for Global System for Mobile Communications (GSM) Subscriber Identity Modules (EAP-SIM),” RFC 4186, January 2006.
- [16] Arkko et al., “Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA),” RFC 4187, January 2006.

- [17] Aboba et al., “PPP EAP TLS Authentication Protocol,” RFC 2716, October 1999.
- [18] Palekar et al., “Protected EAP Protocol (PEAP) Version 2,” Internet Draft, Work in Progress, October 2004.
draft-josefsson-pppext-eap-tls-eap-10.txt
- [19] Funk et al., “EAP Tunneled TLS Authentication Protocol Version 1 (EAP-TTLSv1),” Internet Draft, Work in Progress, March 2006. **draft-funk-eap-ttls-v1-01.txt**
- [20] Cam-Winget et al., “The Flexible Authentication via Secure Tunneling Extensible Authentication Protocol Method (EAP-FAST),” Internet Draft, Work in Progress, January 2007.
draft-cam-winget-eap-fast-06.txt
- [21] Zeilenga K., “Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map,” RFC 4510, June 2006.
- [22] Zeilenga K., “Lightweight Directory Access Protocol (LDAP): Directory Information Models,” RFC 4512, June 2006.
- [23] Harrison R., “Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms,” RFC 4513, June 2006.
- [24] Rosenberg et al., “SIP: Session Initiation Protocol,” RFC 3261, June 2002.
- [25] Cameron, “The Laws of Identity,” May 2005.
- [26] OASIS, “Security Assertion Markup Language 2.0,” March 2005.
- [27] Dory Leifer, “Visitor Networks,” *The Internet Protocol Journal*, Volume 5, No. 3, September 2002.
- [28] William Stallings, “SSL: Foundation for Web Security,” *The Internet Protocol Journal*, Volume 1, No. 1, June 1998.

SEAN CONVERY is CTO at Identity Engines, a venture-backed startup developing innovative identity management solutions for enterprise networks. Prior to Identity Engines, Sean (CCIE® no. 4232) worked for seven years at Cisco Systems, most recently in the office of the security CTO. Sean is best known as the principal architect of the SAFE Blueprint from Cisco and the author of *Network Security Architectures* (Cisco Press, 2004). Sean has presented to or consulted with thousands of enterprise customers around the world on designing secure networks. Before Cisco, Sean held various positions in IT and security consulting during his 14 years in networking. E-mail: **sconvery@idengines.com**

IPv6 Network Mobility

by Carlos J. Bernardos, Ignacio Soto, and María Calderón, Universidad Carlos III de Madrid

The *Internet Protocol* (IP) is currently accelerating the integration of voice and data communications. The Mobile IP protocol enables host mobility support, but several scenarios exist today, such as the provision of Internet access from mobile platforms (for example, planes, trains, cars, etc.), making it necessary to also support the mobility of complete networks. In response to this demand, the *Internet Engineering Task Force* (IETF) has developed the *Network Mobility* (NEMO) *Basic Support Protocol*^[1], enabling IPv6 network mobility.

This article explains the Network Mobility Basic Support Protocol, by first providing a general overview and then examining the details.

Why Network Mobility?

Accelerated by the success of cellular technologies, mobility has changed the way people communicate. As Internet access becomes more and more ubiquitous, demands for mobility are not restricted to single terminals anymore. It is also needed to support the movement of a complete network that changes its point of attachment to the fixed infrastructure, maintaining the sessions of every device of the network: what is known as *network mobility* in IP networks. In this scenario, the mobile network has at least a (mobile) router that connects to the fixed infrastructure, and the devices of the mobile network connect to the exterior through this mobile router.

Support of the roaming of networks that move as a whole is required in order to enable the transparent provision of Internet access in mobile platforms, such as the following:

- *Public transportation systems*: These systems would let passengers in trains, planes, ships, etc. access the Internet from terminals onboard (for example, laptops, cellular phones, *Personal Digital Assistants* [PDAs], and so on) through a mobile router located at the transport vehicle that connects to the fixed infrastructure.
- *Personal networks*: Electronic devices carried by people, such as PDAs, photo cameras, etc. would connect through a cellular phone acting as the mobile router of the personal network.
- *Vehicular scenarios*: Future cars will benefit from having Internet connectivity, not only to enhance safety (for example, by using sensors that could control multiple aspects of the vehicle operation, interacting with the environment and communicating with the Internet), but also to provide personal communication, entertainment, and Internet-based services to passengers.

However, IP networks were not designed for mobile environments. In both IPv4^[2] and IPv6^[3, 4], IP addresses play two different roles. On the one hand, they are *locators* that specify, based on a routing system, how to reach the node that is using that address. The routing system keeps information about how to reach different sets of addresses that have a common network prefix. This address aggregation in the routing system satisfies scalability requirements. On the other hand, IP addresses are also part of the *endpoint identifiers* of a communication, and upper layers use the identifiers of the peers of a communication to identify them. For example, the *Transmission Control Protocol* (TCP), which is used to support most of the Internet applications, uses the IP address as part of the TCP connection identifier.

This dual role played by IP addresses imposes some restrictions on mobility, because when a terminal moves from one network (IP subnet) to another, we would like to *maintain* the IP address of the node that moves (associated to one of its network interfaces) in order not to change the identifier that upper layers are using in their ongoing sessions. However, we also would like to *change* the IP address to make it topologically correct in the new location of the terminal, allowing in this way the routing system to reach the terminal.

Protocols such as the *Dynamic Host Configuration Protocol* (DHCP)^[5, 6] facilitated the portability of terminals by enabling the dynamic acquisition of IP configuration information without involving manual intervention. However, this automation is not enough to achieve real and transparent mobility because it requires the restarting of ongoing transport sessions after the point of attachment changes. The IETF has studied the problem of terminal mobility in IP networks for a long time, and IP-layer solutions exist for both IPv4 (Mobile IPv4^[7, 8]) and IPv6 (Mobile IPv6^[9]) that enable the movement of terminals without stopping their ongoing sessions.

If we focus on IPv6^[3] networks, Mobile IPv6 does not support, as it is now defined, the movement of complete networks. One way of achieving the transparent mobility of all the nodes of a network moving together (for example, in a plane) could be enabling host mobility support in all of them, so they independently manage their mobility. However, this approach has the following drawbacks:

- Host mobility support (for example Mobile IP^[7, 8, 9]) is required in *all* the nodes of the network. This support might not be possible, for example, because of the limited capacities of the nodes (such as in sensors or embedded devices) or because it is not possible to update the software in some older devices. By having a single entity (the mobile router) that manages the mobility of the complete network, nodes of the network do not require any special mobility software to benefit from the transparent mobility support provided by the (mobile) router.

- The signaling exchanged because of the roaming of the network is limited to a single node sending only one message (avoiding “storms” of signaling messages every time the network moves).
- Nodes of the network must be able to attach to the access technology available to connect to the Internet. This requirement might mean that all the nodes of the network should have *Universal Mobile Telecommunications Service* (UMTS) or WiMAX interfaces, for example. On the other hand, by putting this requirement on a single node (the mobile router), nodes of the network can gain access to the Internet through the mobile router, using cheaper and widely available access technologies (for example, *wireless LAN* [WLAN] or Bluetooth).

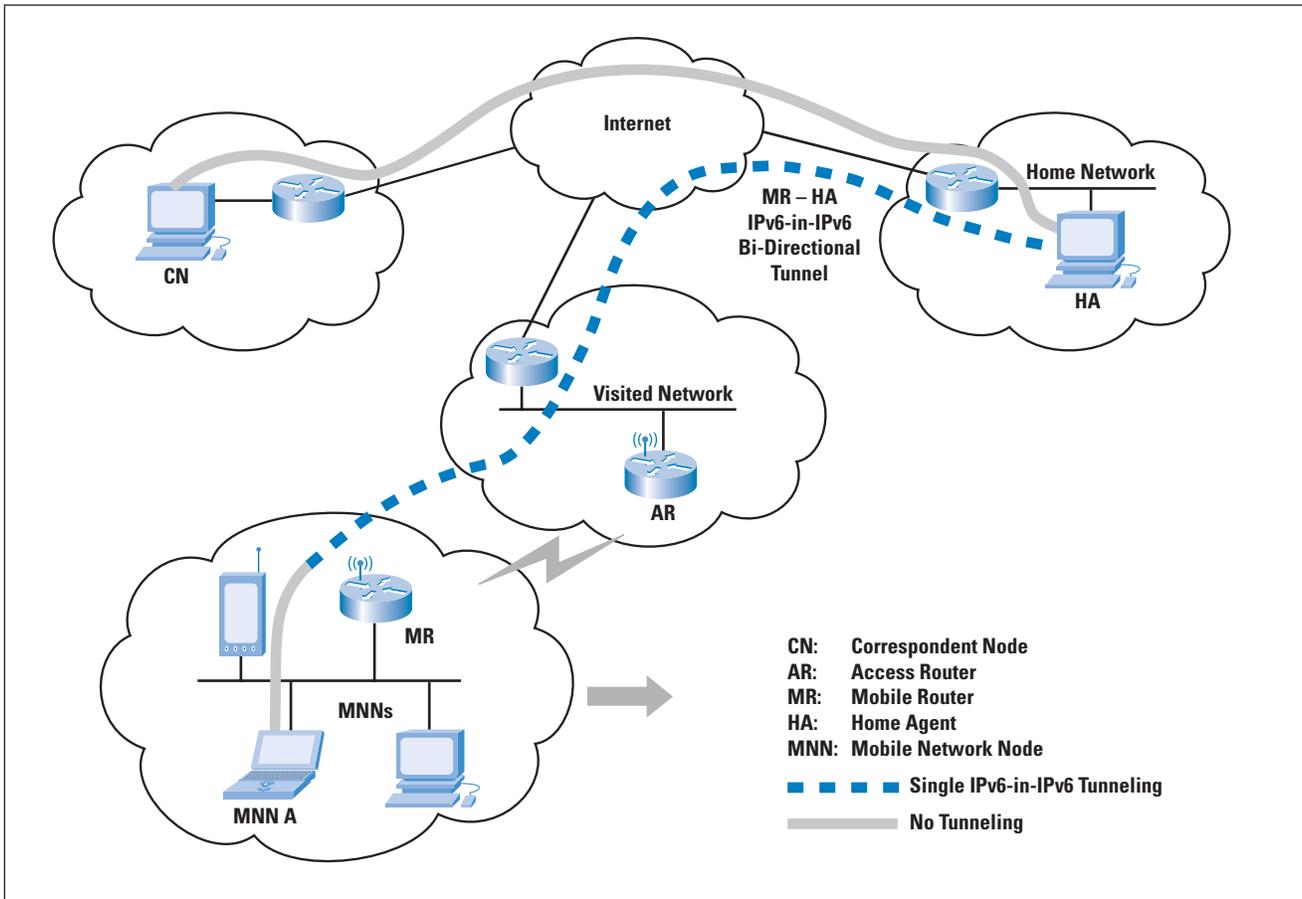
Because of these problems, the IETF *NEMO Working Group* was created to standardize a solution enabling network mobility at the IPv6 layer. The current solution, called the Network Mobility Basic Support Protocol, is defined in RFC 3963^[1].

Operation of the NEMO Basic Support Protocol

A mobile network (known also as a “network that moves,” or *NEMO*) is defined as a network whose attachment point to the Internet varies with time. Figure 1 depicts an example of a network-mobility scenario. The router within the NEMO that connects to the Internet is called the *Mobile Router* (MR). It is assumed that the NEMO is assigned to a particular network, known as its *Home Network*, where it resides when it is not moving. Because the NEMO is part of the home network, the mobile network has configured addresses belonging to one or more address blocks assigned to the home network: the *Mobile Network Prefixes* (MNPs). These addresses remain assigned to the NEMO when it is away from home. Of course, these addresses have topological meaning only when the NEMO is at home. When the NEMO is away from home, packets addressed to the nodes of the NEMO, known as *Mobile Network Nodes* (MNNs), are still routed to the home network. Additionally, when the NEMO is away from home, the mobile router acquires an address from the visited network, called the *Care-of Address* (CoA), where the routing architecture can deliver packets without additional mechanisms.

When any node located at the Internet, known as a *Correspondent Node* (CN), exchanges IP datagrams with a *Mobile Network Node* (MNN; A in Figure 1), the following operations are involved in the communication:

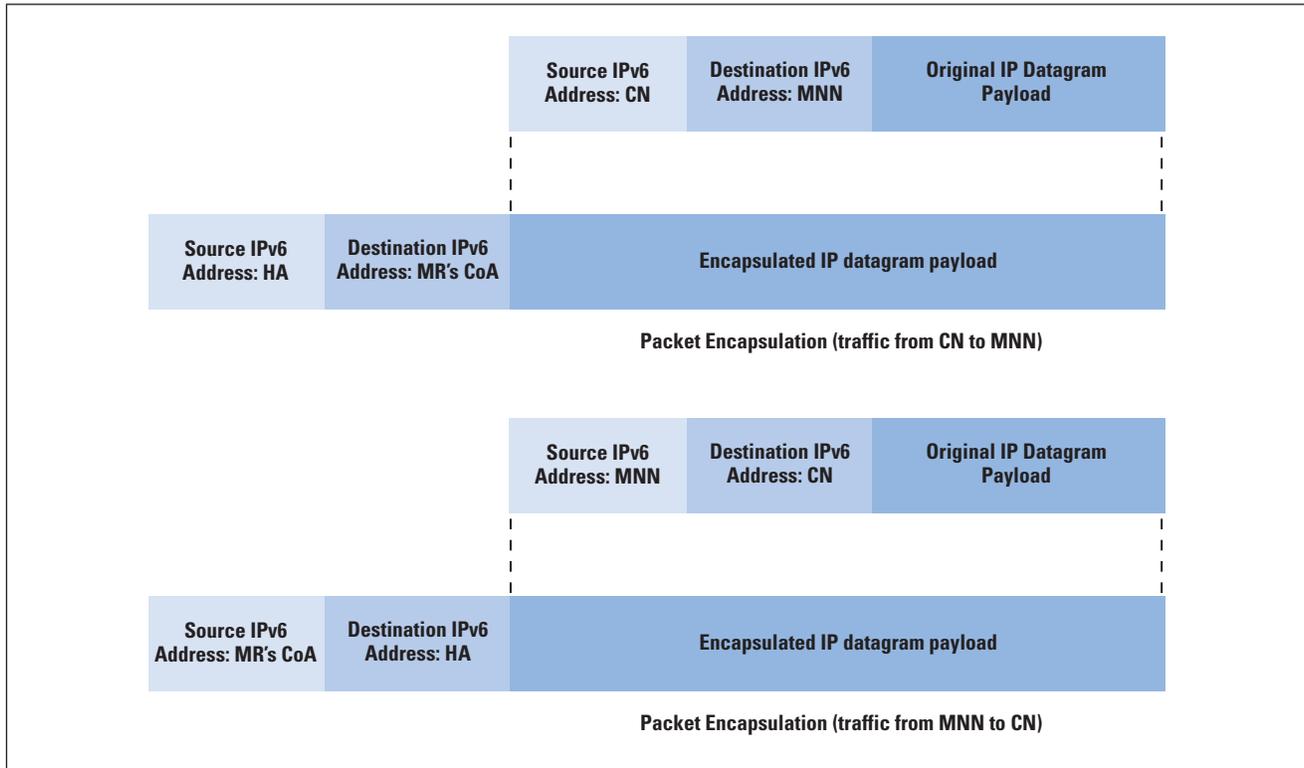
Figure 1: Example of NEMO Basic Support Protocol Operation



1. The correspondent node transmits an IP datagram destined for MNN A. This datagram carries as its destination address the IPv6 address of MNN A, which belongs to the MNP of the NEMO.
2. This IP datagram is routed to the home network of the NEMO, where it is encapsulated inside a new IP datagram by a special node located on the home network of the NEMO, called the *Home Agent* (HA). The new datagram is sent to the CoA of the mobile router, with the IP address of the home agent as source address. This encapsulation (as shown in Figure 2) preserves mobility transparency (that is, neither MNN A nor the correspondent node are aware of the mobility of the NEMO) while maintaining the established Internet connections of the MNN.
3. The mobile router receives the encapsulated IP datagram, removes the outer IPv6 header, and delivers the original datagram to MNN A.

4. In the opposite direction, the operation is analogous. The mobile router encapsulates the IP datagrams sent by MNN A toward its home agent, which then forwards the original datagram toward its destination (that is, the correspondent node). This encapsulation is required to avoid problems with ingress filtering, because many routers implement security policies that do not allow the forwarding of packets that have a source address that appears topologically incorrect.

Figure 2: Overview of NEMO Basic Support Protocol Encapsulation

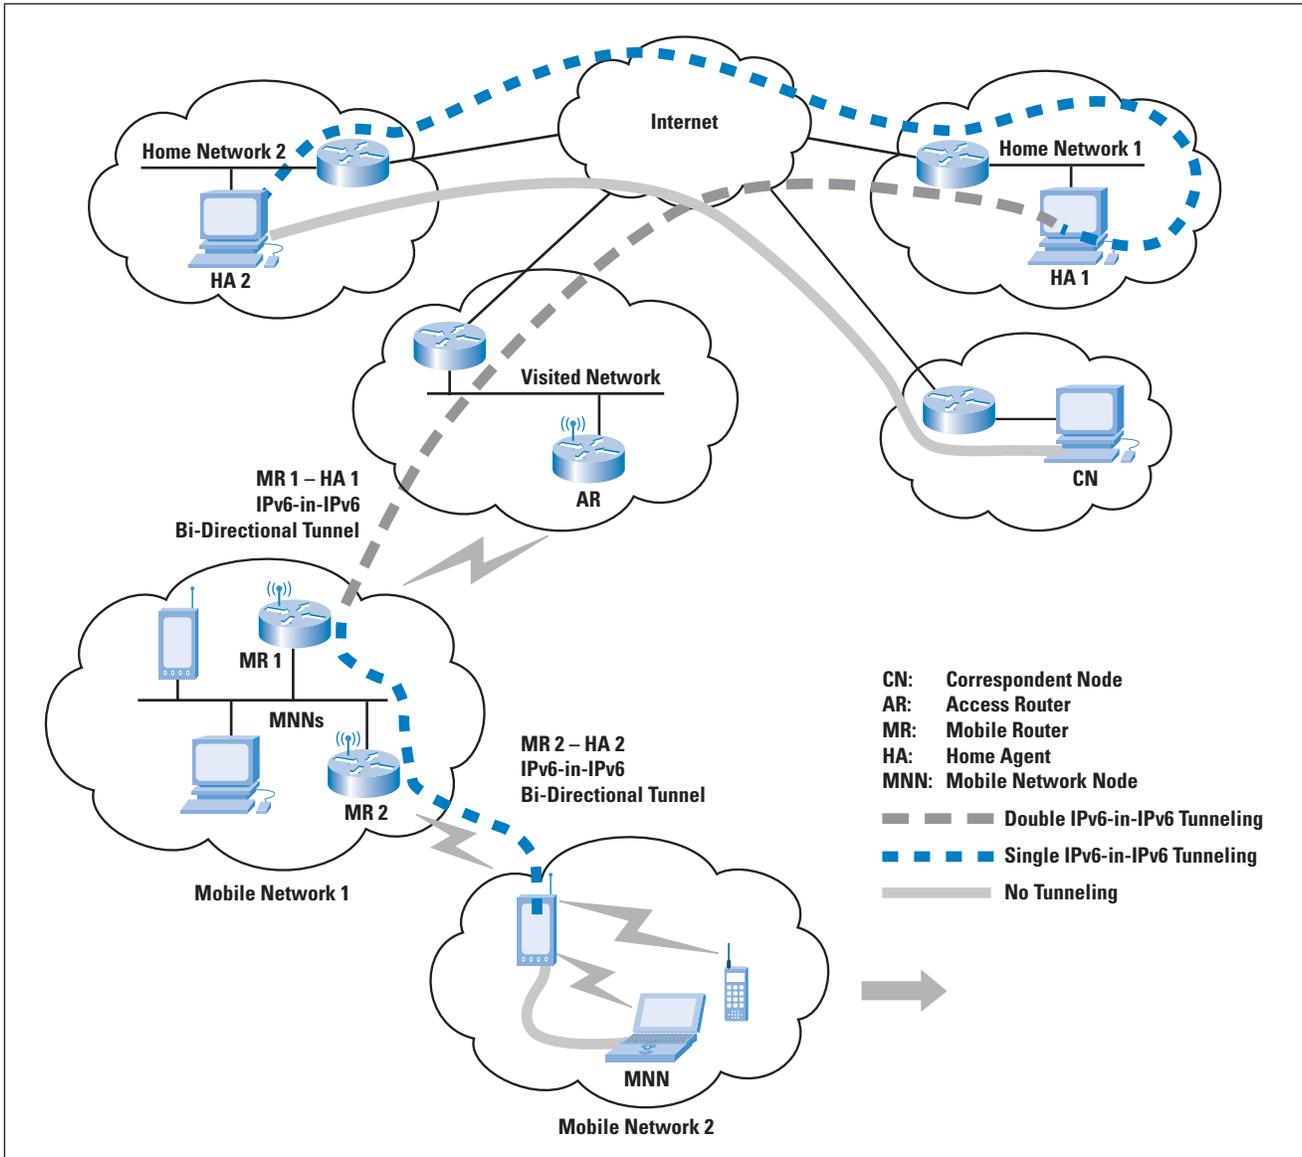


Following are different types of MNNs:

- *Local Fixed Node (LFN)*: This node has no mobility-specific software and therefore cannot change its point of attachment while maintaining ongoing sessions. Its IPv6 address is taken from a MNP of the NEMO to which it is attached.
- *Local Mobile Node (LMN)*: This node implements the Mobile IPv6 protocol; its home network is located in the mobile network. Its *home address* (HoA) is taken from an MNP.
- *Visiting Mobile Node (VMN)*: This node implements the Mobile IP protocol (and therefore, it can change its point of attachment while maintaining ongoing sessions), has its home network outside the mobile network, and it is visiting the mobile network. A VMN that is temporarily attached to a mobile subnet (used as a foreign link) obtains an address on that subnet (that is, its CoA is taken from an MNP).

Additionally, mobile networks can be *nested*. A mobile network is said to be nested when it attaches to another mobile network and obtains connectivity through it (refer to Figure 3). An example is a user who enters a vehicle with his personal area network (mobile network 2) and connects, through a mobile router—like a Wi-Fi enabled PDA—to the network of the car (mobile network 1), which is connected to the fixed infrastructure.

Figure 3: Nested Mobile Network: Operation of the NEMO Basic Support Protocol (multiangular routing)



Protocol Details: NEMO Versus Mobile IPv6

The NEMO Basic Support Protocol is an extension of the solution proposed for host mobility support, *Mobile IPv6* (MIPv6)^[9].

In Mobile IPv6, three mechanisms support the mobility of a host: movement detection, location registration, and traffic tunneling. The NEMO Basic Support Protocol extends some of these mechanisms to support the movement of complete networks. These mechanisms are described next, with those parts that are different from the Mobile IPv6 protocol highlighted.

Movement Detection

In Mobile IPv6, the host needs to discover its own movement, so it can proceed with the required signaling and operations that allow its transparent mobility. Mobile IPv6 defines a generic movement-detection mechanism based on the *Neighbor Discovery Protocol*^[10], which basically consists of listening to *Router Advertisements* (RAs). Routers send these router-advertisement messages, both periodically and in response to a *Router Solicitation* message issued by a host. By looking at the information contained in the router advertisements, a host can determine whether or not it has moved to a new link.

The NEMO Basic Support Protocol does not introduce any change on the movement-detection mechanisms that a mobile router can use.

Location Registration

When a host moves to a new network, it has to configure a new IPv6 address on the visited link (belonging to the IPv6 address space of that visited network): the CoA, and inform the home agent of the movement. In Mobile IPv6, the mobile node (that is, a mobile host) informs its home agent of its current CoA using a mobility message called the *Binding Update* (BU). This message is carried in an IPv6 datagram using a special extension header defined by Mobile IPv6 to encapsulate all messaging related to the creation and management of mobility bindings, called the *mobility header*. The binding-update message contains information required by the home agent to create a mobility binding, such as the home address of the *Mobile Node* (MN) and its CoA, where the home agent should encapsulate all the traffic destined to the mobile node. The home agent replies to the mobile node by returning a *Binding Acknowledgement* (BA) message.

The NEMO Basic Support Protocol extends the binding-update message to convey the following additional information:

- *Mobile Router Flag* (R): The mobile router flag is set to indicate to the home agent that the binding update is from a mobile router. A mobile router can behave as a mobile host: by setting this flag to 0, the home agent does not forward packets destined for the mobile network to the mobile router, but forwards only those packets destined to the home address of the mobile router.

- *Mobile Network Prefix Option*: This option is in the binding update to indicate the prefix information for the mobile network to the home agent. There could be multiple mobile network prefix options if the mobile router has more than one IPv6 prefix in the mobile network and wants the home agent to forward packets for each of these prefixes to the current location of the mobile router.

When the NEMO Basic Support Protocol is used to provide mobility to a complete network, only one binding-update or binding-acknowledgement signaling messages exchange is performed, whereas if the Mobile IP protocol were used by all the nodes of an N -node network, $N \times$ (Binding-update or Binding-acknowledgement) signaling messages synchronized exchanges would be required—usually referred to as a “binding-update signaling storm.”

Mobile IPv6 defines a route-optimization mechanism that enables direct path communication between the mobile node and a correspondent node (avoiding traversal of the home agent). This route optimization is achieved by allowing the mobile node to send binding-update messages also to the correspondent nodes. In this way the correspondent node is also aware of the CoA, where the home address of the mobile node is currently reachable. A special mechanism—called the *Return Routability* (RR) procedure—is defined to prove that the mobile node has been assigned (that is, “owns”) both the home address and the CoA at a particular moment in time^[11], and therefore provides the correspondent node with some security guarantees.

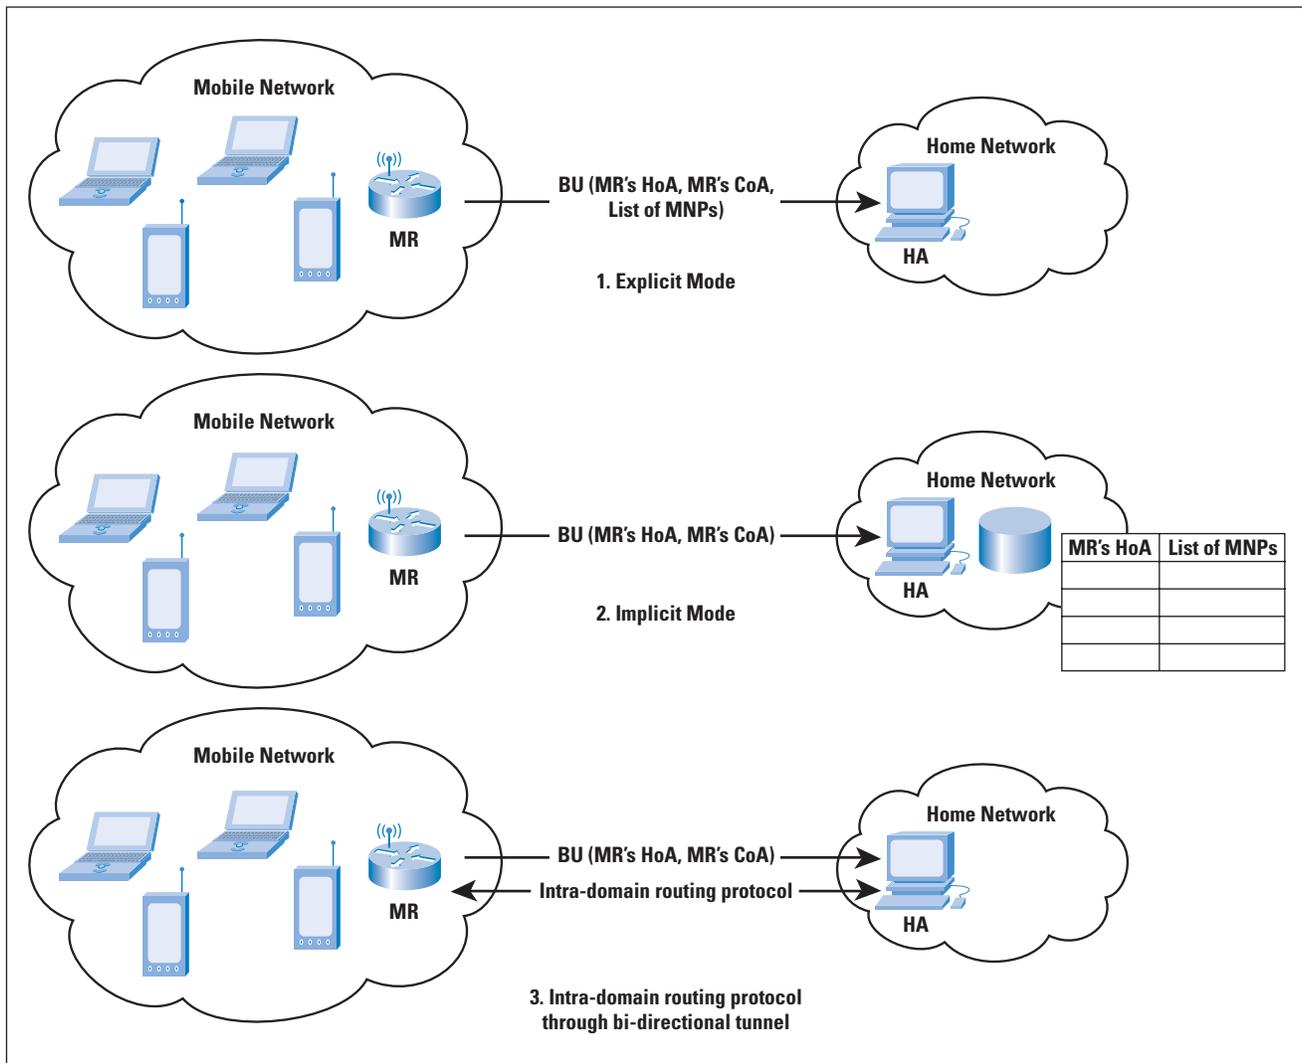
Because of the nature of the network-mobility scenario, the task of providing mobile networks with route-optimization support becomes more complex. The IETF is currently working on this topic^[12, 13, 14].

Traffic Tunneling

In Mobile IPv6, after the mobile node has successfully registered its current location, the home agent starts encapsulating the data traffic destined to the mobile node toward its CoA.

In a NEMO scenario, the home agent forwards not only those IP datagrams arriving at the home network that are destined to the home address of the mobile router, but also all the traffic addressed to any of the mobile-network prefixes managed by the mobile router. The home agent can determine which prefixes belong to the mobile router in three different ways (refer to Figure 4):

Figure 4: NEMO Basic Support Modes of Operation



- *Explicit mode:* The mobile router includes one or more mobile network prefix options in the binding-update message that it sends to the home agent. These options contain information about the mobile-network prefix(es) configured on the mobile network.
- *Implicit mode:* The mobile router does not include prefix information in the binding-update message it sends to the home agent. The home agent determines the mobile-network prefix(es) owned by the mobile router by using any other mechanism (the NEMO Basic Support Protocol does not define any, leaving this prefix determination open to be implementation-specific).

One example would be manual configuration at the home agent mapping the home address of the mobile router to the information required for setting up forwarding for the mobile network.

- *Intradomain Dynamic Routing Protocol through the bidirectional tunnel*: Alternatively to the previous two modes of operation, the home agent and the mobile router can run an intradomain routing protocol (for example, *Routing Information Protocol next generation* [RIPng] or *Open Shortest Path First* [OSPF]) through the bidirectional tunnel. The mobile router can continue running the same routing protocol that it ran when attached to the home link.

Fragmentation may be needed to forward packets through the tunnel between the mobile router and the home agent. In this case, the other end of the tunnel (the home agent of the mobile router) must reassemble the packet before forwarding it to the final destination. This requirement does not contradict the fact that *intermediate* IPv6 routers do not fragment (as opposed to IPv4), because the mobile router and home agent are the actual *ends* of the tunnel.

Performance of the NEMO Basic Support Protocol

The NEMO Basic Support Protocol relies on the creation of a bidirectional tunnel between the mobile router and the home agent to provide transparent mobility support to a complete network. The use of this tunnel causes an additional overhead of 40 bytes per packet, because of the extra IPv6 header added by the encapsulation. The effect of this overhead might be relevant for applications that generate small packets, such as *voice-over-IP* (VoIP) packets, because the 40-byte added overhead may be even bigger than the actual VoIP payload.

The end of the bidirectional tunnel at the side of the mobile router needs to be updated each time the mobile network moves (and also periodically to refresh the binding at the home agent), to reflect the current location of the mobile router. This updating is achieved by the binding-update or binding-acknowledgement signaling exchange between the mobile router and the home agent. As stated previously, only one exchange (two packets, one in each direction) is required per movement, regardless of the number of MNNs that are attached to the mobile router—one of the main advantages of using the NEMO Basic Support Protocol on the mobile router instead of Mobile IPv6 on every node of the mobile network, because the signaling generated by a complete moving network (composed of numerous nodes) is the same as the one generated by a single moving node.

Conclusions

The NEMO Basic Support Protocol^[1] extends the functions of Mobile IPv6 to support the mobility of complete networks. The current specification supports basic mobility, and the IETF is currently working on new enhancements and extensions to provide route-optimization support, multihoming capabilities, and IPv4 support.

Some implementations of the NEMO Basic Support Protocol are already available. For example, the latest Cisco IOS® Software releases provide network mobility support. Open-source implementations also exist, such as the *NEMO Platform for Linux* (NEPL) (<http://www.mobile-ipv6.org/>) and SHISA (<http://www.mobileip.jp/>), for Linux and *Berkeley Software Distribution* (BSD) operating systems, respectively.

References

- [1] Vijay Devarapalli, Ryuji Wakikawa, Alexandru Petrescu, and Pascal Thubert, “Network Mobility (NEMO) Basic Support Protocol,” RFC 3963, January 2005.
- [2] Jon Postel, “Internet Protocol,” RFC 791, September 1981.
- [3] Iljitsch van Beijnum, “IPv6 Internals,” *The Internet Protocol Journal*, Volume 9, No. 3, September 2006.
- [4] Stephen E. Deering and Robert M. Hinden, “Internet Protocol, Version 6 (IPv6) Specification,” RFC 2460, December 1998.
- [5] Ralph Droms, “Dynamic Host Configuration Protocol,” RFC 2131, March 1997.
- [6] Ralph Droms, Jim Bound, Bernie Volz, Ted Lemon, Charles E. Perkins, and Mike Carney, “Dynamic Host Configuration Protocol for IPv6 (DHCPv6),” RFC 3315, July 2003.
- [7] William Stallings, “Mobile IP,” *The Internet Protocol Journal*, Volume 4, No. 2, June 2001.
- [8] Charles E. Perkins, “IP Mobility Support for IPv4,” RFC 3344, August 2002.
- [9] David B. Johnson, Charles E. Perkins, and Jari Arkko, “Mobility Support in IPv6,” RFC 3775, June 2004.
- [10] Thomas Narten, Erik Nordmark, and William A. Simpson, “Neighbor Discovery for IP Version 6 (IPv6),” RFC 2461, December 1998.
- [11] Pekka Nikander, Jari Arkko, Tuomas Aura, Gabriel Montenegro, and Erik Nordmark, “Mobile IP Version 6 Route Optimization Security Design Background,” RFC 4225, December 2005.
- [12] Chan-Wah Ng, Pascal Thubert, Masafumi Watari, and Fan Zhao, “Network Mobility Route Optimization Problem Statement,” Internet Draft, Work in Progress, September 2006.
draft-ietf-nemo-ro-problem-statement-03.txt

- [13] Chan-Wah Ng, Fan Zhao, Masafumi Watari, and Pascal Thubert, “Network Mobility Route Optimization Solution Space Analysis,” Internet Draft, Work in Progress, September 2006. **draft-ietf-nemo-ro-space-analysis-03.txt**
- [14] María Calderón, Carlos J. Bernardos, Marcelo Bagnulo, Ignacio Soto, and Antonio de la Oliva, “Design and Experimental Evaluation of a Route Optimisation Solution for NEMO,” *IEEE Journal on Selected Areas in Communications (J-SAC)*, Issue on Mobile Routers and Network Mobility, Volume 24, Number 9, pages 1702–1716, September 2006.

CARLOS J. BERNARDOS received a telecommunication engineering degree in 2003, and a Ph.D. in telematics in 2006, both from University Carlos III of Madrid. His Ph.D. thesis focused on Route Optimisation for Mobile Networks in IPv6 Heterogeneous Environments. He has been working as a research and teaching assistant in Telematics Engineering since 2003. His current work focuses on IP-based mobile communication protocols. E-mail: **cjbc@it.uc3m.es**

IGNACIO SOTO received a telecommunication engineering degree in 1993, and a Ph.D. in telecommunications in 2000, both from the University of Vigo, Spain. He was a research and teaching assistant in telematics engineering at the University of Valladolid from 1993 to 1999. In 1999 he joined University Carlos III of Madrid, where he has been an associate professor since 2001. His research activities focus on mobility support in packet networks and heterogeneous wireless access networks. E-mail: **isoto@it.uc3m.es**

MARÍA CALDERÓN is an associate professor at the Telematics Engineering Department of University Carlos III of Madrid. She received a computer science engineering degree in 1991 and a Ph.D. degree in computer science in 1996, both from the Technical University of Madrid. She has published more than 20 papers in the fields of advanced communications, reliable multicast protocols, programmable networks, and IPv6 mobility. E-mail: **maria@it.uc3m.es**

More ROAP: Routing and Addressing at IETF68

by Geoff Huston, APNIC

Over the past year or so we have seen a heightened level of interest in Internet routing and addressing. Speculation regarding the future role of the Internet raises the possibility of the Internet supporting as many as hundreds of billions of chattering devices. What does such a future imply in terms of the core technologies of the Internet? Consideration of this topic has prompted a critical examination of the architecture of the Internet, including the scaling properties of routing systems, the forms of interdependence between addressing plans and routing, and the roles of addresses within the architecture.

The March 2007 meeting of the IETF, IETF68, saw some further steps in analysing these topics, and many sessions addressed aspects of routing and addressing. This article reports on these sessions, and includes some conjecture as to what lies ahead.

Plenary ROAP – The Plenary Session on Routing and Addressing

The plenary session presented an overview of the topic, looking at the previous initiatives in routing and addressing, as well as providing some perspectives on the current status of work in this area. There are concerns that the technology platform cannot scale by further orders of magnitude without some changes. Also of concern are the scalability of routing, the “transparency” of the network, renumbering questions, provider-based addressing, and service and traffic engineering and routing capabilities—and these concerns are potentially even more relevant and challenging for tomorrow’s Internet.

Our routing technology does not localize the external effects of local configuration choices. Far from being a protocol that damps instability, the *Border Gateway Protocol* (BGP) is a highly effective amplifier of noise components of routing events. So although it is a remarkably useful information-dissemination protocol, the properties of BGP in an ever-more connected world with ever-finer granularity of information raise some questions about its scaling properties. Will the imposed “noise” of the behaviour of the protocol completely swamp the underlying information content? Will we need to deploy disproportionately larger routers to support a larger network? The prospect here is that routing may become far less efficient because as we simultaneously increase the degree of interconnection and the information load, the inability to effectively localize information creates a far greater load on network routing.

In addition to these observations about routing, there is the continuing suspicion that the semantic load of addresses in the Internet architecture, where an address simultaneously conveys the concepts of “who,” “where,” and “how,” contributes to routing load.

To what extent the semantic intent of endpoint identity (or “id”) can be separated from the semantic intent of network location and forwarding lookup token (or “loc”) is a question of considerable interest. Although the current IP address semantics removes the need to support an explicit mapping operation between identity and location, the cost lies in the inability to support an address plan that is cleanly aligned to network topology, and the inability to cleanly support functions associated with device or network mobility. In the end it is the routing system that carries the consequent load. The questions in this area include an evaluation of the extent to which identity can be separated from location, and the effect of such a measure on the operation of applications. How much of today’s Internet architecture would be affected by such a change, and what would be the resultant benefits if this measure were deployed? Are we necessarily looking at a single model of such an id/loc split, or should we think about this scenario in a more general manner with numerous potential id/loc splits?

Obviously this study of routing and addressing, and the related aspects of name space attributes and mapping and binding properties, has a very broad scope. The larger question posed here is whether we can defer resolution of this problem to a comfortably distant future, or whether its effect on the present network is imminent. Are we accelerating toward some form of near-term technical limit that will cause a significant disruptive event within the deployed Internet, and will volume-based networks economics hold or will bigger networks start to experience disproportionate cost bloat—or worse? Is it time to be alarmed?

The unallocated IPv4 address pool will certainly be exhausted in the coming years, but this sense of alarm over routing and addressing is more about whether there are real limits in the near future in the capability to continue to route the Internet within the deployed platform, using the current technologies, and working within current cost-performance relationships irrespective of whether the addresses in the packet headers are 32 or 128 bits in size. There was a strong sense of “Don’t panic!” in the plenary presentation, with the relatively confident expectation that BGP will be able to carry the routing load of the Internet over the next 3 to 5 years without the need for major protocol “surgery,” and that Moore’s Law will continue to ensure that the capacity and speed of hardware will track the anticipated growth rates. Expectations are that the current technologies and cost-performance parameters will continue to prevail in this time frame.

The *Internet Engineering Steering Group* (IESG) has followed the *Internet Architecture Board’s* (IAB’s) initiative and has begun working with a focus group, the *Routing and Addressing Problem Directorate* (ROAP), to refine the broad space into many more specific work areas, and has assumed a role of coordination and communication across the related IETF activities.

In addition, because a relatively significant research agenda is posed by such long-term questions, the *Routing Research Group* of the *Internet Research Task Force* (IRTF) has been rechartered and, judging by the participation at its most recent meeting, effectively reinvigorated to investigate various approaches to routing that take us well beyond tweaking the existing routing toolset.

Internet ROAP – The Internet Area Meeting

The Internet Area meeting concentrated on aspects of this approach of supporting an identifier/locator split within the architecture of the Internet, and gathering some understanding as to whether this approach would assist with routing scaling. One of the important considerations in this area is working through what could be called boundary conditions of the study. For example, is this matter purely one for protocol stacks within an endpoint, or should distributed approaches that have active elements within the network also be considered? To what extent should a study consider mobility, traffic engineering, *Network Address Translation* (NAT), and *Maximum Transmission Unit* (MTU) behaviour? What appears to be clear at the outset is that this network is not a “clean-slate” network, and any approach should be deployable on the existing infrastructure, should use capability negotiation to trigger behaviours so that deployment can be incremental and piecemeal, should allow existing applications and their identity referential models to operate with no changes, and, hopefully, should have a direct benefit to those parties who decide to deploy the technology.

From the routing perspective, the overall desire is to reduce the growth rates of the interdomain routing space. The desired intent is to reduce the amount of information associated with locators so that locators reflect primarily network topology in such a way that the locators can be efficiently aggregated within the routing system that attempts to maintain a highly stable view of the network topology.

More detailed consideration of the implications of disambiguating aspects of identity from those of network location involves many dimensions—including the structure of the spaces—the mapping functions, and the practicalities of any form of deployment of such a technology.

A critical topic appears to be how an identity-mapping function relates to the forwarding-mapping function. Assuming that the existing name spaces remain unaltered, then the resultant framework appears to require distinct “name-to-identifier” and “identifier-to-locator” mappings and a “locator-to-forwarding” mapping. Where these mapping functions should be performed, who should perform them, when they should be performed, the duration of the validity of the outcomes, whether the mapping function outcomes are relative or universal, the scope and level of granularity in time and space of the map elements, the security of these mapping functions, and whether there is a simple operation in each mapping function or multiple operations all remain undefined at this point.

Other questions include whether the mapping is explicit or implicit, what evidence of a previous mapping operation is held in a packet in a visible manner, and what is occluded from further inspection after the mapping operation has been performed. In addition, what level of state is required in each host, and is there true end-to-end transparency—at what level?

It is likely, at least at this stage of the study, that such a split can have a variety of approaches, both in the intended roles of identifier and location tokens and in their binding. The expectation at this stage of the study is that further ideas will surface, and such ideas will be helpful rather than distracting. It is unclear if a single solution can emerge from this activity, or whether different actors have a sufficiently different set of relative priorities that multiple approaches—each of which expresses different prioritization of functions—are viable longer-term outcomes.

The critical consideration here is that it is unlikely that scaling routing over the longer term to a much larger network is simply a matter of just changing the operation of the routing system itself. Real improvement in this area appears to also require an understanding of the meaning of the objects, or “addresses,” that are being passed within the routing system. The motivation for opening up the identifier or locator space within the Internet area appears to be strongly tied to the notion that if you can unburden some of the roles of the addresses used in routing, and treat these routed tokens as unadorned network locality tokens, then you can gain some additional capability in routing.

Routing ROAP – The Routing Area Meeting

The first part of the Routing ROAP session looked at the trends in the routing system over 2005 and 2006. The overall trend appears to be a system that is increasingly densely interconnected, carrying more information elements, each of which expresses finer levels of granularity in reachability. There appears to be two forms of dynamic BGP load: the BGP “supernova” that burst with an intense BGP update load over some weeks and then disappear, and “background radiation” generators that appear to be unstable at a steady update rate for months or even the entire year.

In looking at scaling the BGP routing environment, one response is that of behavioural changes in local instances of BGP that reduce the potential for unnecessary updates to be propagated beyond a “need-to-know-now” radius. Another response is to consider changes to BGP in terms of additional attributes to BGP updates—such as a “withdrawal-at-origin” flag, or selective advertisement of “next best path”—both of which are intended to limit the span of advertised intermediate transitions while the BGP distance vector algorithm converges to a stable state.

It appears that we could improve our understanding of the operational profile of the routing space, looking particularly at the various forms of pathological routing behaviours and comparing these behaviours against the observations of known control points. Such a study may also lead to some more effective models of projections of the size of the routing space in the near- and medium-term future, and allow some level of quantification as to what “scaling of the routing space” actually implies.

The second part of the Routing ROAP session considered the current status of the routing world, updating some of the observations made at the IAB Routing Workshop and outlining some further perspectives on this space. One critical perspective on BGP is the behaviour of BGP under load. It was noted that most BGP implementations use adaptive responses to peer load, so that BGP attempts to ensure that its peer receives only the most current state information when the peer signals that it is not keeping pace with the update rate.

Another critical factor is the nature of “convergence” in BGP. The claim was made that this problem was the biggest, yet least important, problem with BGP. Convergence delays can be mitigated by *Graceful Restart*, *Nonstop Routing*, and *Fast Reroute*. One of the measures that exacerbates convergence is the use of *Route Reflectors*. The model of information hiding or Route Reflectors is intended to reduce the number of BGP peer sessions and the update load, but the benefits they do achieve are at the cost of slower convergence with a higher message rate during the intermediate-state transitions. Perhaps it is appropriate to consider small-scale changes to BGP behaviour to mitigate the transient BGP update bursts caused by path hunting, including those already mentioned of “withdrawal-at-origin” notification and propagation of backup paths.

The approach advocated here is based on the perspective that BGP is not in danger of imminent collapse, and there is still considerable “headroom” for BGP operation in today’s Internet.

More ROAP?

The routing space is a classic example of the commons, where each party can use routing to solve a multitude of business problems. This includes, for example, using routing to perform load balancing of traffic over a set of transit providers, using a “spot market” in Internet transit services, creating differentiated transit offerings using more specific routes and selective advertisements. The ultimate cost of these local efforts in optimising local business outcomes lies in the increasing bloat in the routing system and the consequent escalation in costs across the entire network in supporting the routing system. There is no way to impose administrative controls on the global routing system, nor have we been able to devise an economic model of routing where the incremental costs of local routing decisions are visible to the originator as true economic costs for the business, and the benefit of a conservative and prudent use of the routing system reaps economic dividends in terms of relatively lower costs for the business.

Like the commons, there are no effective feedback mechanisms to impose constraint on actors in the routing space. Also, like the commons, there is the distinct risk that the cumulative effect of local actions in routing creates a situation that pushes the routing system, either as a whole or in various locales, into a nonfunctioning state.

Whether it needs a sense of urgency to motivate the work, or a sense that there can and should be a better way to plan a future than crude crisis management, the underlying observation is that the routing and address world is fundamental to tomorrow's Internet. Unless we make a concerted effort to understand the various interdependencies and feedback systems that exist in the current environment, and understand the interdependences that exist between network behaviours and routing and addressing models, then I'm afraid that the true potential of the Internet will always lie within our vision—but frustratingly just beyond our grasp.

Further Reading

Following are references to further material on this topic, as presented at IETF68:

- <http://tools.ietf.org/html/draft-iab-raws-report-01>
- http://submission.apricot.net/chatter07/slides/future_of_routing/apia-future-routing-john-scudder.pdf
- http://submission.apricot.net/chatter07/slides/future_of_routing/apia-future-routing-jari-arkko.pdf
- <http://www3.ietf.org/proceedings/07mar/slides/plenaryw-3.pdf>
- <http://www3.ietf.org/proceedings/07mar/agenda/intarea.txt>
- <http://www3.ietf.org/proceedings/07mar/agenda/rtgarea.txt>
- <http://www1.tools.ietf.org/group/irtf/trac/wiki/RRG>
- <http://www.ietf.org/IESG/content/radir.html>

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. The author of numerous Internet-related books, he is currently the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of the Internet Society from 1992 until 2001. E-mail: gih@apnic.net

Opinion: Is It Time to Replace SMTP?

by Dave Crocker, Brandenburg Internet Working

The first Internet (ARPANET) e-mail, sent 35 years ago, was remarkably similar to a basic text e-mail of today: From, To, CC, Subject, Date, followed by lines of text, and the familiar @-sign in addresses. The right side of the address changed from a simple string into the multilevel domain name that we now use. The body can now be a set of multimedia attachments rather than just lines of text, but it can still be in its original, simpler form. The means of moving mail was the *File Transfer Protocol* (FTP) in the early 1970s. The current mechanism, the *Simple Mail Transfer Protocol* (SMTP)^[1a, 1b], was not created until 10 years later, but a mere 25 years of use is not bad, either.

All of the technical specifications for e-mail have undergone many changes over the years, but a core requirement has been to protect the installed base of users and operators by incrementally adding features as options, rather than by performing wholesale replacement of any infrastructure service component. E-mail has changed the way we communicate, yet it is also now viewed as having a serious problem: As the Internet grew, it acquired the full mixture of participants, some of whom do not make nice neighbors.

Frustration with the effect of abusive users is often expressed as a belief that the solution lies in replacing some or all of the core technology of the e-mail service, or even by moving to an entirely different paradigm, such as querying Webpages using *Really Simple Syndication* (RSS)^[2]. Although different paradigms make sense for some forms of human communications, what is forgotten in these pleas for massive change is the power of the classic mail model, whether by paper or by electrons: Spontaneous or occasional communication requires the ability to “push” the message to the recipient, without prior arrangement. This ability is, of course, also what leaves the door open for abuse—anyone may walk in, uninvited and unwanted.

The alternative proposals might work well enough for ongoing, regular communication among people who already know each other. And for most of us, that is probably 80 percent of our exchanges, or more. Unfortunately, as soon as anyone starts worrying about the remaining 20 percent, these alternative approaches require cascading hacks, producing a design that looks no better than what we have today, except that it is based strictly on theory rather than decades of practice. It is easy for a paper proposal to beat a deployed system; making it work as promised is, of course, more difficult.

Mantra

I have developed a simple mantra, in response to calls for replacing today's Internet mail:

0. The basic problems we are experiencing with e-mail are really based on undesirable social behaviors, long popular outside the Internet. The Internet enables broader reach, to more victims, and in much shorter time spans, but the core misbehaviors have existed for all of recorded human history. We should not assume that there are technical solutions to social problems.
1. The beginning of changing a human service is to gain community consensus about the change that is needed, because a mechanism will not be successful unless it is perceived as needed. Only then can the engineers work on designing the change.
2. When there is community consensus about the way that e-mail needs to be changed, the folks who are currently contributing to its 35-year evolution need to try to find a way to add the desired features to the existing service. Given the record of accomplishment of e-mail, the odds seem favorable that any new requirement can also be satisfied without disrupting the installed base.
3. When that effort fails, it will be time to create a replacement infrastructure.

Alas, as those who track e-mail abuse technical discussions are well aware, we have not completed Step 1. As soon as we try to formulate community consensus about basic messaging communication policies, discussion devolves into cacophony or marginalized community fragments. It is certain that there will eventually be a change required for e-mail, which we cannot fit into the current service, but we do not yet have any evidence that e-mail abuse is going to produce that requirement.

Trust Models

One hopeful sign is that we do have a solid set of efforts to evolve e-mail to support mechanisms that are based on trust. This evolution begins with the ability to associate a validated identity to a message and then requires assessing the "safety" of that identity's owner. Until recently, only the IP address of the last-hop sending SMTP server could be used as an identifier. Using addresses as identifiers sounds reasonable at first glance, but turns out to have long-term scaling and administrative problems. As a result, there has been a broad effort to find ways to use domain names, which are more stable, and they align better with organizational boundaries. This process is well under way, with the recent IETF standardization of the *Domain Keys Identified Mail* (DKIM)^[3] message-signing specification, as well as path-based registration schemes, such as Sender-ID^[4] and SPF^[5].

That took about 5 years. And now comes the hard part: developing a range of *assessment mechanisms*—sometimes generically called *reputation services*—that satisfy requirements for quality, strength, convenience, and stability. Assessment services tell recipients whether the author of the message, or the service that sent it, can be trusted. Some mechanisms need to work for small groups, others need to work for mass-market business-to-consumer mailings, and others need to work among business partners. A few startup companies have recently joined the few, surviving volunteer services, to satisfy this need. It is too early to tell whether they will suffice, or whether additional services will be needed. What is important is that these services are generally regarded as producing good results.

For the long term it seems likely that this capability will result in an Internet mail service that is logically split into two types of traffic. One has substantial trust associated with its messages, so that they can be delivered with a reasonable degree of comfort. The other is the current, open-to-all service that requires heavy filtering and the use of various heuristics, to reduce the effect of abuse mail. If the first traffic flow is sufficiently successful, filters for the second can become much more stringent. The aggregate effects of these changes will be that wanted mail is likely to be received and identified much more reliably, and unwanted mail is more likely to be rejected.^[8, 9]

So the current Internet mail technical infrastructure is safe, right? Well, maybe.

Enhancements?

What gets less attention, but perhaps should worry us more, is the general lack of user-level functional enhancement for e-mail. What users can do with e-mail, today, is pretty much the same as they could do 25 years ago. The evolution of Internet mail has been primarily in support of performance, reliability, and scaling. Although important, they have not produced functional changes that are apparent to end users. Human communication is a very rich space, yet most e-mail is limited to a narrow range of styles: person-to-person informal communications, and informal, unstructured group communications. Toss in some very basic, one-way “transactional” mail, such as order confirmations from businesses to their customers, and that about covers it.

Instead, new functions for human collaboration have tended to appear in new services. *Instant Messaging* (IM), blogging, and wikis are the most popular examples. In each case, they rely on a centralized service, rather than the highly distributed model that e-mail uses. Users must all go to a single, centralized address to obtain a given service. Most of the IM world does not even know that there are two (!) Internet standards for distributed IM—*Extensible Messaging and Presence Protocol* (XMPP)^[6] and SIMPLE^[7]. Even for these standards, most of their production use tends to be within noninteroperable, centralized services.

Is there something about e-mail that is a barrier to functional enhancements for end users?

For these new services, the interservice relaying that is at the core of e-mail is absent. Indeed, centralized services are easier to create and operate than are distributed services, but they also carry scaling, administration, and control challenges. So the issue is not so much what is easier, but who will do the work—and when? With a centralized service, all the interesting work is done by the single provider. For a distributed model, like e-mail, the work is shared across participating organizations. The Internet was designed to avoid single points of failure (and failure), so it is ironic that these new services risk exactly these problems.

For a distributed model, like e-mail, to add end-user functions, useful adoption is required by all user software that participates, and possibly by all the intermediate, relaying services. The adoption is in three parts: agreeing on the enhancement, modifying existing software, and making it available to users. These are daunting barriers, so the appeal of centralized services is clear: a single organization decides what to change, changes it, and makes it available to end users with, at most, a natural software upgrade.

Interorganization partnerships provide the best argument for adoption of distributed services, because they do not naturally permit agreement on a central point of control. The counterforce is, again, the simplification (for the partners) that comes from agreeing to use independent third-party services. The scaling problem here is with end users having to juggle a large number of independent services. Note the emergence of IM clients that support a variety of independent IM services.

Perhaps the real danger to e-mail is not its wholesale and traumatic replacement, stemming from frustration about abuses, but a gradual attrition, as portions of its traffic move to services that evolve more quickly, but leave end users with a complicated array of narrow, specialized, and noninteroperable venues.

References

- [1a] Postel, J. B., “Simple Mail Transfer Protocol,” RFC 821, August 1982.
- [1b] Klensin, J., “Simple Mail Transfer Protocol,” RFC 2821, April 2001.
- [2] Really Simple Syndication Specifications,
<http://www.rss-specifications.com/rss-specifications.htm>

- [3] Allman, E., et al., “DomainKeys Identified Mail (DKIM) Signatures,” February 2007. (*RFC publication pending.*)
<http://dkim.org/specs/draft-ietf-dkim-base-10.html>
- [4] Lyon, J. and Wong, M., “Sender ID: Authenticating E-Mail,” RFC 4406, April 2006.
- [5] Wong, M. and Schlitt, W., “Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1,” RFC 4408, April 2006.
- [6] Saint-Andre, P. (ed.), “Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence,” RFC 3921, October 2004.
- [7] Campbell, B. (ed.), Rosenberg, J., Schulzrinne, H., Huitema, C., and Gurle, D., “Session Initiation Protocol (SIP) Extension for Instant Messaging,” RFC 3428, December 2002.
- [8] Crocker, D., “Challenges in Anti-Spam Efforts,” *The Internet Protocol Journal*, Volume 8, No. 4, December 2005.
- [9] Klensin, J., “Taking Another Look at the Spam Problem,” *The Internet Protocol Journal*, Volume 8, No. 4, December 2005.

DAVE CROCKER is a principal with Brandenburg InternetWorking. He has authored or contributed to most Internet mail standards, and an assortment of e-mail products and businesses, as well as working on facsimile, security, e-commerce, and EDI. He received the 2004 *IEEE Internet Award* for his work on e-mail. Dave is a contributor to the development efforts for DKIM, CSV, and BATV, motivated by a strong desire to protect more than 30 years of professional investment that is being threatened by spamming. E-mail: dcrocker@bbiw.net

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

ARIN Board Advises Internet Community on Migration to IPv6

The *American Registry for Internet Numbers* (ARIN) and the other *Regional Internet Registries* (RIRs) have distributed Internet Protocol version 6, IPv6, alongside IPv4 since 1999. To date, ARIN has issued both protocol versions in tandem and has not advocated one over the other. ARIN has closely monitored trends in demand and distribution for both protocol versions with the understanding that the IPv4 available resource pool would continue to diminish.

The available IPv4 resource pool has now been reduced to the point that ARIN is compelled to advise the Internet community that migration to IPv6 is necessary for any applications that require ongoing availability from ARIN of contiguous IP number resources. On 7 May 2007, the ARIN Board of Trustees passed the following resolution:

“Whereas, community access to *Internet Protocol* (IP) numbering resources has proved essential to the successful growth of the Internet; and,

Whereas, ongoing community access to *Internet Protocol version 4* (IPv4) numbering resources can not be assured indefinitely; and,

Whereas, *Internet Protocol version 6* (IPv6) numbering resources are available and suitable for many Internet applications,

Be it Resolved, that this Board of Trustees hereby advises the Internet community that migration to IPv6 numbering resources is necessary for any applications which require ongoing availability from ARIN of contiguous IP numbering resources; and,

Be it Ordered, that this Board of Trustees hereby directs ARIN staff to take any and all measures necessary to assure veracity of applications to ARIN for IPv4 numbering resources; and,

Be it Resolved, that this Board of Trustees hereby requests the ARIN Advisory Council to consider Internet Numbering Resource Policy changes advisable to encourage migration to IPv6 numbering resources where possible.”

Implementation of this resolution will include both internal and external components. Internally, ARIN will review its resource request procedures and continue to provide policy experience reports to the Advisory Council. Externally, ARIN will send progress announcements to the ARIN community as well as the wider technical audience, government agencies, and media outlets. ARIN will produce new documentation, from basic introductory fact sheets to FAQs on how this resolution will affect users in the region. ARIN will focus on IPv6 in many of its general outreach activities, such as speaking engagements, trade shows, and technical community meetings. For more information, visit ARIN’s IPv6 Information Center at:

<http://www.arin.net/v6/v6-info.html>

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

Editorial Advisory Board

Dr. Vint Cerf, VP and Chief Internet Evangelist
Google Inc, USA

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

David Farber
Distinguished Career Professor of Computer Science and Public Policy
Carnegie Mellon University, USA

Peter Löthberg, Network Architect
Stupi AB, Sweden

Dr. Jun Murai, General Chair Person, WIDE Project
Vice-President, Keio University
Professor, Faculty of Environmental Information
Keio University, Japan

Dr. Deepinder Sidhu, Professor, Computer Science &
Electrical Engineering, University of Maryland, Baltimore County
Director, Maryland Center for Telecommunications Research, USA

Pindar Wong, Chairman and President
Verifi Limited, Hong Kong

*The Internet Protocol Journal is published quarterly by the Chief Technology Office, Cisco Systems, Inc.
www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com*

Copyright © 2007 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website are the property of their respective owners.

Printed in the USA on recycled paper.



The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive, M/S SJ-7/3
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

PRSRT STD U.S. Postage PAID PERMIT No. 5187 SAN JOSE, CA
--