The Internet Protocol Journal

September 2010 Volume 13, Number 3

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

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

ISSN 1944-1134

From The Editor

Technology advances—such as improvements in display technology, battery life, processor capabilities, and communications systems—have all contributed to making *mobile devices* the most important area for Internet growth. In order to fully support these devices, the IETF developed *Mobile IP* many years ago, and it has continued to work on the general area of IP mobility. We have covered some of this work in previous issues of IPJ, and this time we look at *Proxy Mobile IPv6* (PMIPv6), which is being standardized by the IETF. The article is by Ignacio Soto, Carlos J. Bernardos, María Calderón, and Telemaco Melia.

Deployment of IPv6 is progressing, albeit slowly. In several upcoming articles we will examine some transition technologies or implementation details that can make this deployment easier, and above all, transparent, to the end user. In our first article, Dan Wing and Andrew Yourtchenko explain the concept of “Happy Eyeballs” as applied to dual-stack IPv4/IPv6 systems.

Domain Name System Security Extensions (DNSSEC) have recently been applied to the Internet system of root servers. For details, see our “Fragments” section, where you will also find a statement from the *Number Resource Organization* (NRO) regarding the results of a recent IPv6 readiness study.

Once again, please remember to check your subscription expiration date and take the necessary steps if you wish to continue receiving this journal. It’s not too late to renew and get back on the distribution list, even if your subscription expired some time ago. You can find your subscription ID and expiration date either on the back page of your copy or on the envelope that it came in. In order to access your record, click the “Subscriber Services” link on our webpage at `www.cisco.com/ipj` and enter your e-mail address and the subscription ID. The system will send you a link that allows direct access to your record, and you can update your address and renew your subscription. If you no longer have access to the e-mail you used when you subscribed or have forgotten your subscription ID, just send a message to `ipj@cisco.com` and we will make the necessary changes for you.

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

PMIPv6: A Network-Based Localized Mobility Management Solution

by Ignacio Soto, Universidad Politécnica de Madrid; Carlos J. Bernardos, and María Calderón, Universidad Carlos III de Madrid; and Telemaco Melia, Alcatel Lucent Bell Labs

Traditional IP mobility procedures[4] are based on functions residing in both the mobile terminal and the network. Recently, we have been assisting in a shift in IP mobility protocol design, mostly focusing on solutions that relocate mobility procedures from the mobile device to network components. This new approach, known as *Network-Based Localized Mobility Management* (NetLMM), allows conventional IP devices (for example, devices running standard protocol stacks) to roam freely across wireless stations belonging to the same local domain. This property is appealing from the operator's viewpoint because it allows service providers to enable mobility support without imposing requirements on the terminal side (for example, software and related configuration). For this purpose the *Internet Engineering Task Force* (IETF) has standardized *Proxy Mobile IPv6* (PMIPv6)[1].

This article details the Proxy Mobile IPv6 protocol, providing a general overview and an exhaustive description of a few selected functions.

Why Network-Based Localized Mobility?

The ability to move while being connected to a communication network is very attractive for users, as demonstrated by the success of cellular networks. However, while designing the IP stack, mobility was not retained as a requirement and, as a consequence, IP does not natively support mobility. The reason is a very basic design choice adopted in IP, both in IPv4[2] and in IPv6[3], namely that addresses have two roles: they are used as locators and identifiers at the same time.[16]

IP addresses are *locators* that specify, by means of the routing system, how to reach the node (more properly, the *network interface*) that is using a specific destination address. The routing system keeps information about how to reach different sets of addresses that have a common network prefix, thus improving scalability of the system itself. However, IP addresses are also *identifiers* used by upper-layer protocols (for example, the *Transmission Control Protocol* [TCP]) to identify the endpoints of a communication channel. Additionally, names of nodes are translated by the *Domain Name System* (DNS) to IP addresses (which, in that way, play the role of node identifiers).

The linking of these two roles (*locators* and *identifiers*) is appealing because name resolution of the peer with whom we want to communicate and location finding translate to the same problem (that is, no translation mechanism is needed). However, the negative side effect is that supporting mobility becomes difficult.

Mobility implies separating the identifier role from the location one. From the identification standpoint, the IP address of a node should never change, but from the location point of view the IP address should change each time the node moves, showing its current location within the routing hierarchy (that is, the IP subnet to which the node is currently attached).

The IETF has studied the problem of terminal mobility in IP networks for a long time. It has developed IP-layer solutions for both IPv4 (Mobile IPv4[4], [5]) and IPv6 (Mobile IPv6[6]), enabling the movement of terminals and providing transparent service continuity. These solutions, being IP-based, are independent of the Layer 2 technologies. They provide Mobile Nodes with a permanent address (the *Home Address* [HoA]) to be used as identifier, and a temporal address (the *Care-of Address* [CoA]) to be used as locator. The CoA changes in each IP subnet visited by the Mobile Node. An entity in the network, the *Home Agent,* binds both addresses with the help of signaling generated by the Mobile Node. The Home Agent serving a Mobile Node must be placed in the subnet where the Home Address of that Mobile Node is topologically correct (the home network).

Although Mobile IP enables a host to move (that is, change the point of attachment in an IP network) while keeping session continuity, this ability is not sufficient for true mobility. Enabling efficient handoffs is an additional and critical requirement. Because the IP handoff latency is affected by the time required to exchange signaling between the Mobile Node and the Home Agent, a new family of solutions proposes to use a local Home Agent (that is, a Home Agent closer to the Mobile Node) to provide mobility in a local domain; that is, to provide localized mobility support. Changing the point of attachment within the local domain requires only signaling to the local Home Agent, allowing faster signaling messages exchange because it is limited within the local domain. This approach is attractive because users typically move in localized environments (for example, they commute between their living homes and their work places) that can be covered with localized domains. Examples of these types of solutions are "Regional Registrations for IPv4"[7] or "Hierarchical Mobile IPv6 for IPv6"[8]. Note that the term "localized" refers to a particular area from the point of view of the IP network topology, but depending on the access technology, geographically the area can be large, as happens when applying a localized mobility approach to cellular networks.

A common feature of Mobile IP and the localized mobility proposals mentioned previously is that all of them are *host-based*. Mobile Nodes must signal themselves to the network when their location changes and must update routing states in the Home Agent, in the local Home Agent, or in both. This situation also raises the problem of complex security configurations to authenticate those signaling exchanges and modifications of routing states.

Therefore, the IETF decided to work on a solution for NetLMM[10, 11], compounding the advantages of a network-based approach with the benefits of localized mobility management strategies. In NetLMM the network provides mobility support, although the Mobile Node does not participate in IP mobility procedures. That is, network operators can provide mobility support without requiring additional software and complex security configuration in the Mobile Nodes. Thus the deployment of network-based mobility solutions is greatly facilitated. Moreover, the Mobile Node can implement any global mobility solution, because the localized one is transparent and independent from it.

There are several target scenarios for Network-Based Localized Mobility Management[9]:

- Large campus networks with *Wireless Local-Area Network* (WLAN) access: Users move with IP standard devices (that is, no additional hardware or software is required) within a campus that provides WLAN access and mobility support.
- Advanced beyond-third-generation (3G) networks: Cellular operators have been important promoters in the development of the NetLMM solution in the IETF. *Universal Mobile Telecommunications System* (UMTS) and *General Packet Radio Service* (GPRS) networks use a proprietary network-based localized mobility mechanism to provide mobility support for user data traffic (typically IP). This mechanism is based on the GPRS Tunneling Protocol[11], a special-purpose solution developed for *Third-Generation Partnership Project* (3GPP) networks that uses TCP/IP application layer tunnels. A standardized NetLMM protocol for the Internet has important advantages:
  - Reduced costs in network management and in equipment supporting the technology (because of economy of scale)
  - Easier extension of mobility support to other technologies
  - Easier integration with other networks
- Other more-complex scenarios involving network mobility, as in automotive scenarios[12], could benefit from a NetLMM approach to support mobility.

With these advantages in mind, the IETF has standardized a protocol to provide Network-Based Localized Mobility support in IP networks, the *Proxy Mobile IPv6* (PMIPv6) protocol.

Operation of Proxy Mobile IPv6

The main idea of PMIPv6 is that the mobile node is not involved in any IP layer mobility-related signaling. The Mobile Node is a conventional IP device (that is, it runs the standard protocol stack). The purpose of PMIPv6 is to provide mobility to IP devices without their involvement. This provision is achieved by relocating relevant functions for mobility management from the Mobile Node to the network.

PMIPv6 provides mobility support within a localized area, the *Localized Mobility Domain* (LMD) or PMIPv6 domain. While moving within the LMD, the Mobile Node keeps its IP address, and the network is in charge of tracking its location. PMIPv6 is based on *Mobile IPv6* (MIPv6), reusing the Home Agent concept but defining nodes in the network that must signal the changes in the location of a Mobile Node on its behalf.

Figure 1: Network Entities in Proxy Mobile IPv6

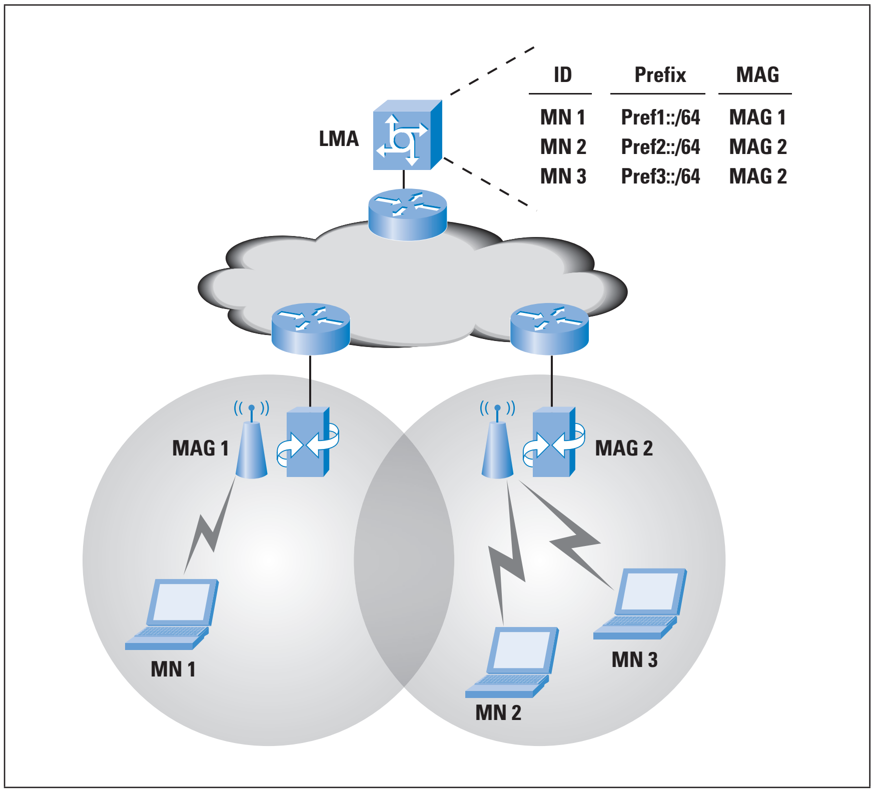

ID	Prefix	MAG
MN 1	Pref1::/64	MAG 1
MN 2	Pref2::/64	MAG 2
MN 3	Pref3::/64	MAG 2

The functional entities in the PMIPv6 network architecture (refer to Figure 1) include the following:

- *Mobile Access Gateway* (MAG): This entity performs the mobility-related signaling on behalf of the Mobile Nodes attached to its access links. The MAG is usually the access router for the Mobile Node, that is, the first-hop router in the Localized Mobility Management infrastructure. It is responsible for tracking the movements of the Mobile Node in the LMD. An LMD has multiple MAGs.

- *Local Mobility Anchor* (LMA): This entity within the core network maintains a collection of routes for each Mobile Node connected to the LMD. The routes point to MAGs managing the links where the Mobile Nodes are currently located. Packets sent or received to or from the Mobile Node are routed through tunnels between the LMA and the corresponding MAG. The LMA is a topological anchor point for the addresses assigned to Mobile Nodes in the LMD, meaning that packets with those addresses as destination are routed to the LMA.

The basic operation of PMIPv6 follows. When a Mobile Node enters a PMIPv6 domain, it attaches to an access link provided by a MAG. The MAG proceeds to identify the Mobile Node, and checks if it is authorized to use the network-based mobility management service. If it is, the MAG performs mobility signaling on behalf of the Mobile Node (see in Figure 2 the signaling when the Mobile Node enters the PMIPv6 domain). The MAG sends to the LMA a *Proxy Binding Update* (PBU) associating its own address with the identity of the Mobile Node (for example, its *Media Access Control* [MAC] address or an identifier related to its authentication in the network). Upon receiving this request, the LMA allocates a prefix to the Mobile Node. Then the LMA sends to the MAG a *Proxy Binding Acknowledgment* (PBA) including the prefix allocated to the Mobile Node. It also creates a *Binding Cache* entry and establishes a bidirectional tunnel to the MAG. The MAG sends *Router Advertisement* messages to the Mobile Node, including the prefix allocated to the Mobile Node, so the Mobile Node can configure an address (stateless autoconfiguration). The Mobile Node can alternatively use stateful address autoconfiguration mechanisms. For simplicity, we assume in the rest of the article that the stateless address autoconfiguration mechanism is used, except when indicated otherwise.

Figure 2: Signaling When a Mobile Node Connects to the PMIPv6 Domain

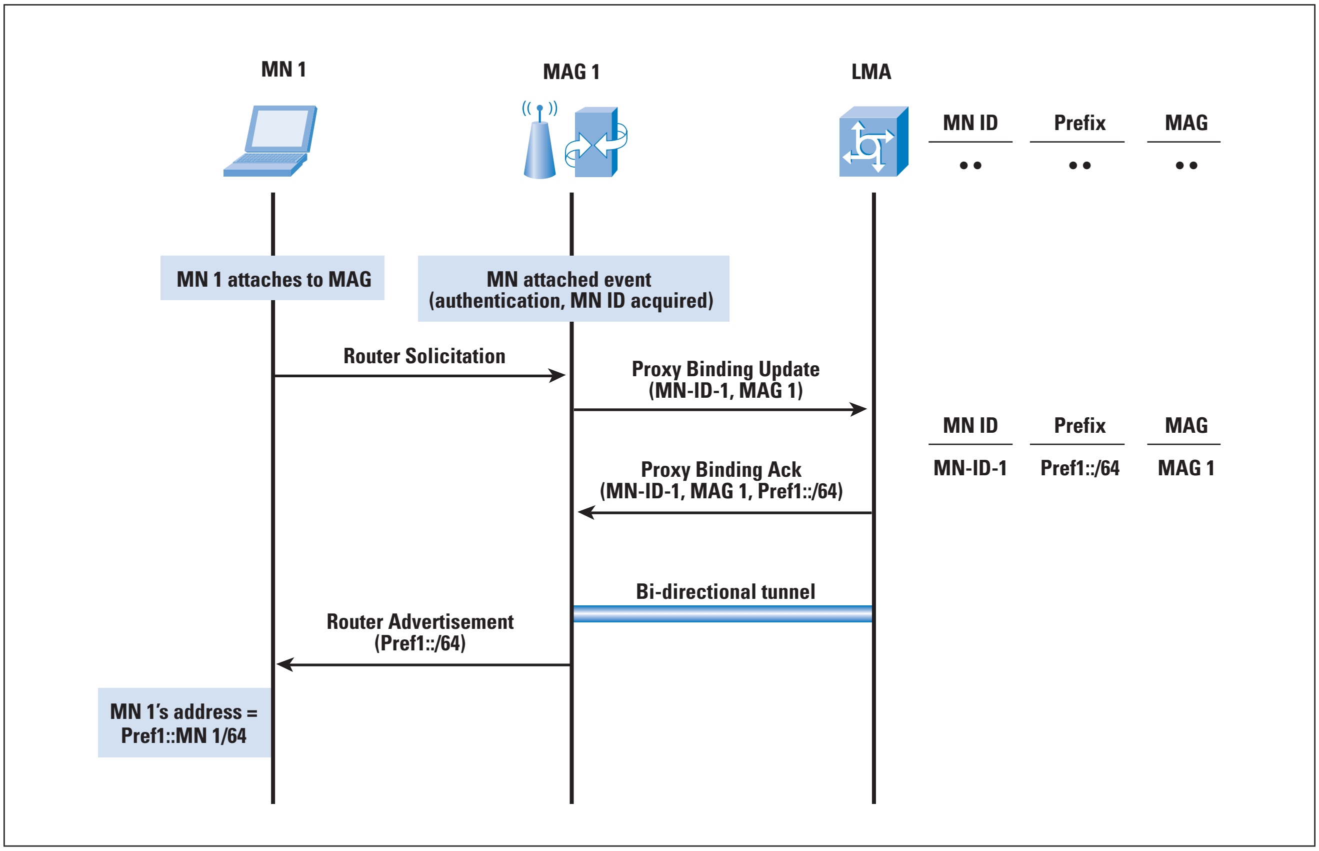
MN 1
MAG 1
LMA
MN ID
Prefix
MAG
••
••
••
MN 1 attaches to MAG
MN attached event (authentication, MN ID acquired)
Router Solicitation
Proxy Binding Update (MN-ID-1, MAG 1)
MN ID
Prefix
MAG
MN-ID-1
Pref1::/64
MAG 1
Proxy Binding Ack (MN-ID-1, MAG 1, Pref1::/64)
Bi-directional tunnel
Router Advertisement (Pref1::/64)
MN 1's address = Pref1::MN 1/64

Whenever the Mobile Node moves, the new MAG updates the location of the Mobile Node in the LMA and advertises the same prefix to the Mobile Node (through Router Advertisement messages), thereby making the IP mobility transparent to the Mobile Node. In this way the Mobile Node keeps the address configured when it first enters the LMD, even after changing its point of attachment within the network, and the LMD appears as a single link from the perspective of the Mobile Node. It should be noted that all the MAGs configure the same link local address for a specific Mobile Node. That is, the Mobile Node will never see a change in its default route configuration.

The bidirectional tunnel between the LMA and the MAG and associated routing states in both LMA and MAG manage the Mobile Node data plane. Downlink packets sent to the Mobile Node from outside of the LMD arrive to the LMA, which forwards them through the tunnel to the serving MAG. The MAG, after decapsulation, sends the packets to the Mobile Node directly through the access link. Uplink packets that originated in the Mobile Node are sent to the LMA from the MAG through the tunnel, and then are forwarded to the destination by the LMA. Traffic originated inside the LMD and directed to a Mobile Node also inside the LMD follows a similar procedure, going through two tunnels from the originating MAG, to the LMA, and then to the destination MAG. It should be noted that PMIPv6 allows a MAG to short-circuit the tunneling in case two mobile nodes directly communicate through any of its interfaces.

Protocol Details

We next describe the PMIPv6 primary functions. Because PMIPv6 is based on the Mobile IPv6 protocol format, we will highlight the differences and extensions to MIPv6. Readers interested in knowing all protocol details should refer to the RFC[1].

Entering a PMIPv6 Domain

The Mobile Node enters the PMIPv6 domain by attaching to an access link. PMIPv6 defines a new functional entity, the MAG, typically residing in the access router. The MAG detects the attachment of the Mobile Node to the access link. The only access link types supported in PMIPv6 are point-to-point links; other types of links can be used as long as they are configured to emulate point-to-point links.

The MAG, upon detecting a Mobile Node attachment, verifies if the Mobile Node is eligible to the network-based mobility management service. Specific procedures to achieve this verification are out of the scope of the PMIPv6 standard. A Mobile Node that uses the mobility support service is identified by the network entities using a *Mobile Node Identifier* (MN-ID). The MN-ID must be stable and unique for the Mobile Node throughout the PMIPv6 domain, but the exact nature of this identifier is not specified. Possible examples are the Mobile Node MAC address or an identifier obtained as part of the Mobile Node authentication procedure.

After the MAG identifies the Mobile Node, authorizes its use of the NetLMM service, and acquires its Mobile Node Identifier, the MAG sends a PBU to the LMA; that is, it sends a registration request on behalf of the Mobile Node to the LMA. The PBU message is based on the MIPv6 *Binding Update* (BU) message with some extensions, but whereas the BU is sent by the Mobile Node, the PBU is sent by the MAG on behalf of the Mobile Node. A flag in the message is used to indicate that it is a PBU and not a BU. The PBU has as source address (and also in the alternate CoA option, if present) the global address configured in the egress interface of the MAG. This address is called *Proxy-CoA* in PMIPv6 terminology and is used by the LMA as locator of the Mobile Node. In the PBU, unlike in the BU, a Home Address destination option is not present; instead a *Mobile Node Identifier Option*[13] has to be included with the Mobile Node Identifier, which is used to identify the Mobile Node throughout the PMIPv6 domain.

The PBU also contains additional information, such as the access link technology, a handoff indicator, the requested lifetime for the registration, and other optional data. The *handoff indicator* is a new mobility option defined in PMIPv6 that allows the MAG to signal the LMA whether the PBU originated upon network attachment or upon handover of a Mobile Node (if known by some unspecified mechanisms), and that information could be useful to support advanced functions such as multihoming. Examples of values of the handoff indicator include: a Mobile Node entering the PMIPv6 domain, a reregistration to update the registration lifetime, a handoff between MAGs, or a handoff between interfaces of the Mobile Node.

Upon sending the PBU, the MAG creates a Binding Update List entry[6] for the Mobile Node. Note that this data structure in Mobile IPv6 is maintained by the Mobile Node to keep track of its bindings, but consequently to the PMIPv6 philosophy, the MAG maintains a *Binding Update List* (BUL) storing the bindings of the Mobile Nodes attached to it. The information in the Binding Update List allows the MAG to link the information about the Mobile Node, the interface in the MAG to which the Mobile Node is connected, and the LMA serving it, among others.

When the LMA receives the PBU sent by the MAG, it first checks that the message is correct according to the PMIPv6 specification, rejecting the registration otherwise. If the LMA accepts the PBU, it has to verify if its *Binding Cache* contains an entry for the Mobile Node identified in the PBU. When a Mobile Node first enters the PMIPv6 domain, the LMA cannot find an entry in its Binding Cache and has to create a new one. The Binding Cache entry is an extended version of the data structure defined for the Binding Cache entries in Mobile IPv6[6].

The entry in the Binding Cache has a flag to indicate that it is a proxy registration, and it links all the information related to the Mobile Node, including its identification and the MAG serving it; that is, the location of the Mobile Node. If there is no entry for the Mobile Node in the Binding Cache (that is, the Mobile Node is entering the PMIPv6 domain), the LMA allocates one or more network prefixes to the Mobile Node. These prefixes are called *Home Network Prefixes,* and it must be noted that at least one network prefix is assigned per Mobile Node.

If the LMA cannot allocate a network prefix to a Mobile Node, it has to reject the registration. The address(es) that the Mobile Node uses while inside the PMIPv6 domain are configured from those Home Network prefixes. The decision of allocating one or more network prefixes depends on a global policy in the PMIPv6 domain or a per-Mobile Node policy. When the registration request is accepted, the LMA creates a *Binding Cache Entry* (BCE) with the accepted values for the registration, including the Mobile Node Identifier, the Proxy CoA (the address of the MAG serving the Mobile Node), and the Home Network prefix(es) allocated to the Mobile Node.

Upon BCE creation, the LMA creates an IPv6-in-IPv6 bidirectional tunnel, if one does not already exist, to the MAG sending the PBU. The LMA sets up forwarding routes through the tunnel for any traffic received that is addressed to the Home Network prefixes of the Mobile Node. Finally, the LMA creates a *Proxy Binding Acknowledgment* (PBA) and sends it to the corresponding MAG. The PBA message is based on the MIPv6 *Binding Acknowledgment* (BA) message with a few more extensions, including a flag that indicates that the message is a Proxy Binding Acknowledgement. The PBA informs the MAG about the registration request result, if it has been rejected (and why, using a status code) or accepted. The PBA contains the Mobile Node Identifier and the Home Network prefixes allocated to the Mobile Node. Unlike the Binding Acknowledgment, the PBA does not include a type 2 routing header (that in the Binding Acknowledgment includes the Home Address of the Mobile Node). Also the PBA is received and processed by the MAG, and not by the Mobile Node.

If the PBA confirms that the registration request has been accepted for the Mobile Node, the MAG creates an IPv6-in-IPv6 bidirectional tunnel, if one does not already exist, to the LMA. The MAG sets up forwarding routes, through the tunnel, for uplink or downlink packets received or sent from or to the Mobile Node. The MAG also updates the Binding Update List entry to reflect the accepted binding registration values.

Upon network attachment and during the PBU or PBA procedure, the Mobile Node can send a *Router Solicitation* in the access link as part of the standard neighbor discovery procedures. The MAG should not reply to this Router Solicitation until the registration in the LMA has been successfully completed. When the MAG receives the PBA indicating a successful registration, the MAG sends a Router Advertisement to the Mobile Node announcing the Home Network prefix(es). The Mobile Node can then apply the stateless address autoconfiguration mechanism or the stateful one (using the *Dynamic Host Configuration Protocol* [DHCP]) according to the indication in the Router Advertisement. For supporting DHCP, a DHCP relay agent has to be present in every MAG in the domain, and the relay agent must include in the link-address field of the *Relay Forward* message an IPv6 address from the Home Network prefix, to indicate to the DHCP server the range of addresses it can assign.

The PMIPv6 specification, as mentioned previously, supports only point-to-point access links with the Mobile Nodes. An interesting use case is to have a broadcast access link and to emulate point-to-point links with the Mobile Nodes to be able to apply the PMIPv6 specification. This case raises the problem of sending Router Advertisements that should be received only by the corresponding Mobile Node, and not by other Mobile Nodes present in the broadcast link. There are several ways to send these advertisements. The Router Advertisements could be sent to the IPv6 link-local address of the Mobile Node that the MAG can learn from the source address of router solicitations sent by the Mobile Node, or by some other unspecified means. Another possibility is to send Router Advertisements to the all-nodes multicast address at the IP layer but to the Link Layer 2 address of the Mobile Node.

Changing MAG in a PMIPv6 Domain

The complete signaling for supporting the change of attachment by a Mobile Node in a PMIPv6 domain is described in Figure 3.

When a Mobile Node leaves a link, the event is detected by the corresponding MAG. The mechanism for Mobile Node movement detection is not specified in PMIPv6, but some possible options are link-layer events or an *IPv6 Neighbor Unreachability Detection* event. The MAG that detects that the Mobile Node has left the link must send a PBU with a Mobile Node de-registration request to the LMA. Upon receiving a PBA replying to the PBU or after a timer, the MAG deletes all the states associated with a specific Mobile Node.

When the LMA receives a PBU with a de-registration request for a Mobile Node with a valid entry in the Binding Cache, it sends the corresponding PBA and starts a timer. During the period defined by the timer the LMA drops any packets received for the Mobile Node. The use of this timer allows the LMA to receive a PBU from a new MAG updating the location of the Mobile Node. If the PBU is not received during that time, the LMA deletes the state associated with the Mobile Node.

Figure 3: Signaling When a Mobile Node Changes Point of Attachment

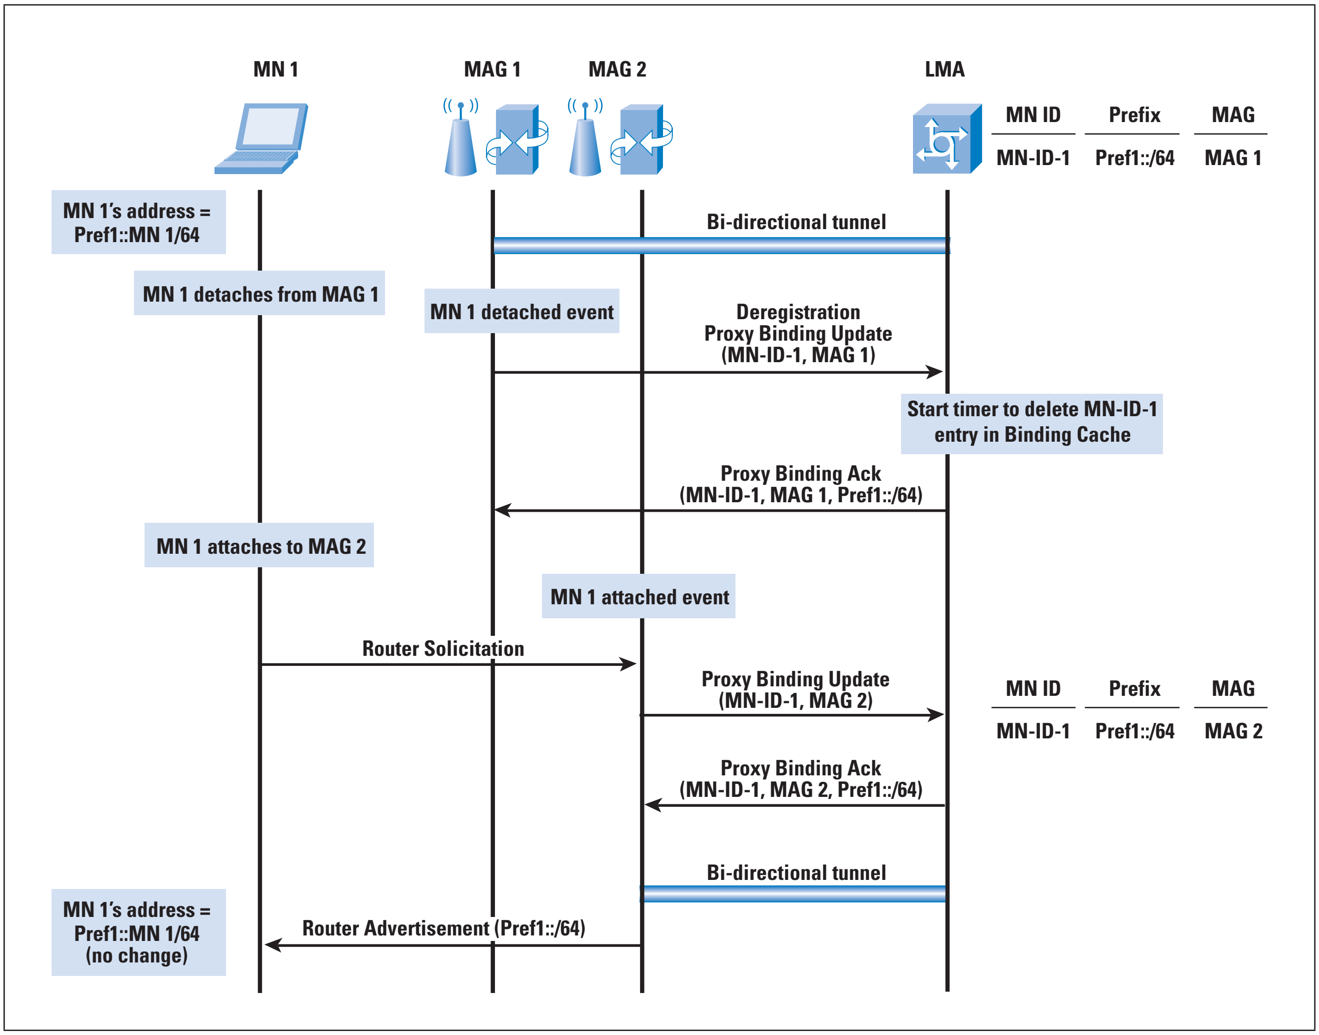
MN 1
MAG 1
MAG 2
LMA
MN ID
Prefix
MAG
MN-ID-1
Pref1::/64
MAG 1
MN 1's address = Pref1::MN 1/64
Bi-directional tunnel
MN 1 detaches from MAG 1
MN 1 detached event
Deregistration Proxy Binding Update (MN-ID-1, MAG 1)
Start timer to delete MN-ID-1 entry in Binding Cache
Proxy Binding Ack (MN-ID-1, MAG 1, Pref1::/64)
MN 1 attaches to MAG 2
MN 1 attached event
Router Solicitation
Proxy Binding Update (MN-ID-1, MAG 2)
MN ID
Prefix
MAG
MN-ID-1
Pref1::/64
MAG 2
Proxy Binding Ack (MN-ID-1, MAG 2, Pref1::/64)
Bi-directional tunnel
MN 1's address = Pref1::MN 1/64 (no change)
Router Advertisement (Pref1::/64)

In a handoff situation the Mobile Node, after leaving a link, attaches to a new access link associated with a new MAG. The new MAG detects the Mobile Node and sends a PBU to the LMA on behalf of the Mobile Node. The LMA receives and processes the PBU, and detects that there is already a Binding Cache entry for that Mobile Node (the same Mobile Node Identifier). The LMA updates the Binding Cache entry with the new information, in particular with the Proxy CoA (egress IPv6 address) of the new MAG, updating also the tunnel and routing information for handling the traffic from or to the Mobile Node. The LMA sends a PBA to the new MAG in which it includes the Home Network prefix(es) already assigned to the Mobile Node. This scenario allows the new MAG to send a Router Advertisement with the same network prefix information as the Mobile Node received from the previous MAG. As stated before, the Mobile Node does not detect a link change and it keeps the same address(es). To make the change of link completely transparent to the Mobile Node, it must also continue receiving the Router Advertisements from the same link-local and link layer address; otherwise the Mobile Node would detect a change of default router. We describe how this problem is addressed in the next section.

Home Network Emulation and Address Uniqueness

MAGs must ensure that Mobile Nodes do not detect link changes when moving in a PMIPv6 domain; that is, MAGs must provide a home network emulation to the Mobile Nodes. To achieve this emulation, all the MAGs in the PMIPv6 domain must send, to a particular Mobile Node, Router Advertisements with the same network prefix information, as described previously. Additionally, the source IPv6 link-local address and the source link layer address in Router Advertisements sent to a Mobile Node must never change, independently of the MAG sending them. Therefore, the PMIPv6 specification requires all the MAGs to use, in any access link to which a particular Mobile Node attaches, the same link-local and link layer address.

PMIPv6 proposes two ways to meet this requirement:

- Configure a fixed link-local and link layer address to be used in all the access links in a PMIPv6 domain.
- Generate at the LMA the link-local address to be used by MAGs with a particular Mobile Node, and send it to the serving MAG through PMIPv6 signaling messages.

Both of these configuration methods are also helpful to guarantee address uniqueness in the access links of the PMIPv6 domain. The global addresses are always unique because all links are point-to-point and only one Mobile Node uses unicast global addresses over that link. Link-local addresses are used by the MAG and the Mobile Node on the link and a collision is possible. However, because the PMIPv6 specification requires that the link-local address used by the different MAGs with a particular Mobile Node is always the same while the Mobile Node moves across the PMIPv6 domain, the collision problem can happen only when the Mobile Node enters the PMIPv6 domain.

When a Mobile Node enters the domain, we must rely on *Duplicate Address Detection* (DAD) to detect a collision. If we use a globally unique link-local address for all the MAGs in the PMIPv6, then it is easy for the MAGs to respond to DAD requests from Mobile Nodes, because MAGs always know the address they must defend. If the link-local address to be used by the MAG with a Mobile Node is generated in the LMA, then it is desirable that the MAG learns that link-local address (that is, completes the PMIPv6 registration procedure) to defend it before the Mobile Node carries out the DAD procedure. You can ensure the MAG can learn this address by ensuring that the Layer 2 attachment is not completed until finishing the PMIPv6 signaling registration, or by configuring the PMIPv6 registration procedure in such a way that it is likely to be completed before the default waiting time of a DAD procedure.

Security Considerations

As with Mobile IPv6 signaling, PMIPv6 signaling is very sensitive to security threats, because it changes routing states of nodes in the network on behalf of the Mobile Nodes. PMIPv6 specification recommends using *IP Security* (IPsec) to protect the signaling exchanges between the MAGs and the LMA. A security association is needed between MAGs and the LMA, but how it is created is not defined. Two cases are possible:

- The network elements (LMA and MAGs) belong to the same operator.
- The elements belong to different operators with an agreement for roaming support.

In both scenarios, creating the security association is an affordable problem.

Traffic Handling in a PMIPv6 Domain

Traffic sent to any address belonging to a Home Network prefix is received by the LMA, the anchor point for those addresses. The LMA forwards the traffic through the tunnel to the MAG serving the Mobile Node, and the MAG decapsulates the packets and forwards them to the Mobile Node through the access link. Packets sent by the Mobile Node are forwarded by the MAG through the tunnel to the LMA. The LMA decapsulates the packets and forwards them to the destination. If a MAG has data traffic that originated in one of its access links and is destined to another of its access links, it can forward the traffic locally to avoid the forwarding through the LMA. This forwarding is done according to a policy configured in the MAG.

Performance Considerations

PMIPv6 presents two performance advantages compared with MIPv6. First, the LMA is a local network entity, so in principle the delay of sending signaling to the LMA will be lower than sending signaling to a remote Home Agent. And second, because the tunnel required to handle the traffic is terminated in the MAG instead of in the Mobile Node (as happens in MIPv6), we avoid the overhead of having a tunnel (two IP headers) over the radio interface. This overhead avoidance is relevant because bandwidth resources are scarcer over the air interface than in the backhaul network.

IPv4 Support Considerations

PMIPv6 acknowledges the existence of a dual-stack mobile host. To this end there are ongoing efforts to standardize IPv4 support for PMIPv6 operations. The extensions defined in [14] specify how to assign an IPv4 Home Address to a mobile host accessing the PMIPv6 domain. That is, the MAG—upon Mobile Node detection attachment and verification that the Mobile Node is eligible for PMIPv6 service—inserts in the PBU an “IPv4 Home Address Request Option.”

The LMA, upon reception of the PBU message, assigns an IPv6 *Home Network Prefix* (HNP) or an IPv4 Home Address by attaching an "IPv4 Home Address Reply Option" to the PBA. How the information is delivered to the Mobile Node depends on the interface between the Mobile Node and the MAG, possible examples being DHCP or *Internet Key Exchange Version 2* (IKEv2). The Mobile Node—independent of the method deployed—configures the HNP and the IPv4 Home address assigned by the LMA, thus supporting both IPv4- and IPv6-based applications.

Conclusions

PMIPv6 is a promising specification that allows network operators to provide localized mobility support without relying on mobility functions or configuration present in the mobile nodes. This reality greatly eases the deployment of the solution.

The IETF is currently working in the *Network-Based Mobility Extensions* (netext) Working Group on extending the PMIPv6 specification to add functions such as enhanced multihoming and intertechnology handoff support, and localized routing for traffic between MAGs to avoid going through the LMA. Additionally, the *Multicast Mobility* (multimob) Working Group is working on the support of multicast in PMIPv6.

References

[1] S. Gundavelli (Ed.), K. Leung, V. Devarapalli, K. Chowdhury, and B. Patil, "Proxy Mobile IPv6," RFC 5213, August 2008.

[2] Jon Postel, "Internet Protocol," RFC 791, September 1981.

[3] Stephen E. Deering and Robert M. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," RFC 2460, December 1998.

[4] William Stallings, "Mobile IP," *The Internet Protocol Journal,* Volume 4, Number 2, June 2001.

[5] Charles E. Perkins, "IP Mobility Support for IPv4," RFC 3344, August 2002.

[6] David B. Johnson, Charles E. Perkins, and Jari Arkko, "Mobility Support in IPv6," RFC 3775, June 2004.

[7] E. Fogelstroem, A. Jonsson, and C. Perkins, "Mobile IPv4 Regional Registration," RFC 4857, June 2007.

[8] H. Soliman, C. Castelluccia, K. ElMalki, and L. Bellier, "Hierarchical Mobile IPv6 (HMIPv6) Mobility Management," RFC 5380, October 2008.

[9] J. Kempf (Ed.), "Problem Statement for Network-Based Localized Mobility Management (NETLMM)," RFC 4830, April 2007.

[10] J. Kempf (Ed.), "Goals for Network-Based Localized Mobility Management (NETLMM)," RFC 4831, April 2007.

[11] 3GPP TS 29.060, "GPRS Tunnelling Protocol (GTP) across the Gn and Gp interface," 2009. Available at: **http://www.3gpp.org/ftp/Specs/html-info/29060.htm**

[12] Ignacio Soto, Carlos J. Bernardos, Maria Calderon, Albert Banchs, and Arturo Azcorra, "NEMO-Enabled Localized Mobility Support for Internet Access in Automotive Scenarios," *IEEE Communications Magazine,* Vol. 47, No. 5, May 2009.

[13] A. Patel, K. Leung, M. Khalil, H. Akhtar, and K. Chowdhury, "Mobile Node Identifier Option for Mobile IPv6 (MIPv6)," RFC 4283, November 2005.

[14] R. Wakikawa and S. Gundavelli, "IPv4 Support for Proxy Mobile IPv6," RFC 5844, May 2010.

[15] Carlos J. Bernardos, Ignacio Soto, and María Calderón, "IPv6 Network Mobility," *The Internet Protocol Journal,* Volume 10, Number 2, June 2007.

[16] Dave Meyer, "The Locator Identifier Separation Protocol (LISP)," *The Internet Protocol Journal,* Volume 11, Number 1, March 2008.

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 was an associate professor from 2001 until 2010. In 2010, he joined Universidad Politécnica de Madrid as associate professor. His research activities focus on mobility support in packet networks and heterogeneous wireless access networks. E-mail: **isoto@dit.upm.es**

CARLOS J. BERNARDOS received a telecommunication engineering degree in 2003, and a Ph.D. in telematics in 2006, both from the University Carlos III of Madrid, where he worked as a research and teaching assistant from 2003 to 2008, and since then as an associate professor. His Ph.D. thesis focused on route optimization for mobile networks in IPv6 heterogeneous environments. He has published more than 30 scientific papers in prestigious international journals and conferences, and he also contributes to the IETF. He served as TPC chair of WEEDEV 2009 and as guest editor of *IEEE Network*. E-mail: **cjbc@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 40 papers in the fields of advanced communications, reliable multicast protocols, programmable networks, and IPv6 mobility. E-mail: **maria@it.uc3m.es**

TELEMACO MELIA received his Informatics Engineering degree in 2002 from the Polytechnic of Turin, Italy, and his Ph.D. in Mobile Communications from the University of Goettingen in April 2007. From June 2002 to December 2007 he worked at NEC Europe Ltd. in Heidelberg, Germany, in the Mobile Internet Group. He worked on IPv6-based Mobile Communication focusing on IP mobility support across heterogeneous networks and resource optimization control. In September 2008 he joined Alcatel Lucent Bell Labs. He is currently working on interworking architectures spanning 3GPP, WiMAX forum, and IETF standardization bodies. His main research interests include wireless networking and next-generation networks. He his author of more than 20 publications and he actively contributes to the IETF. E-mail: **telemaco.melia@alcatel-lucent.com**

Improving User Experiences with IPv6 and SCTP

by Dan Wing and Andrew Yourtchenko, Cisco Systems

To be successful, new technologies must improve the user experience. In the process of finding the best way to deploy a new technology, several approaches are typically conceived, written down, tried, and possibly discarded. This article addresses two such approaches for *Internet Protocol Version 6* (IPv6) and the *Stream Control Transmission Protocol* (SCTP)[10].

Modern web browsers, web servers, and operating systems support IPv4 and IPv6, and several major content providers already support IPv6, including Google, NetFlix, and Facebook. However, their properties are not generally available over IPv6 because of a conflict between IPv6 technology and their business realities.

The technology in web browsers and operating systems involves doing *Domain Name System* (DNS) queries for *AAAA* and *A* resource records and then attempting to connect to the resulting IPv6 and IPv4 addresses *sequentially.* If the IPv6 path is broken (or slow), this connection can take a long time before it falls back to trying IPv4. This process is especially painful on typical websites that retrieve objects from different hosts—each failure incurs a delay. The combination of operating system and web browser results in delays from 20 seconds to several minutes if the IPv6 path is broken[2]. The typical message flow of a TCP client is shown in Figure 1. Clearly, this delay is unacceptable to users. Users avoid this delay by disabling IPv6[3] or avoiding IPv6-enabled websites.

The problem of broken IPv6 networks is relatively widespread[6]. Providing content is a business—either directly (for example, streaming movies) or indirectly (for example, selling advertising). If users suffer delays viewing IPv6-enabled content (because of the technology reasons described previously), they will have an incentive to visit other websites. This scenario means lost revenue and is unacceptable to the business. Considering that all of the customers on today's Internet can reach IPv4 content, it is a business risk to enable IPv6 because some customers will suffer delays attempting to view IPv6 websites. Major content providers have been monitoring the situation and have published results[7] showing that the IPv6 failure rate is too high to enable IPv6 *AAAA* for their content.

IPv6 problems have several causes. It is new technology, and monitoring of IPv6 connectivity is not yet on par with that of IPv4 because of single-point tunnels, unmanaged tunnels[11], accidentally misconfigured firewalls, and router and link failures can more easily cause outages on IPv6. Many applications remain IPv4-only, or network administrators are relying on dual-stack equipment to transparently fail over to IPv4 during IPv6 outages.

However, such failover is never transparent to users—it takes many seconds or minutes! To avoid these problems, the content provider has only one choice: don't provide *AAAA* records if users might experience broken or slow IPv6.

Figure 1: Behavior of a Typical Web Browser

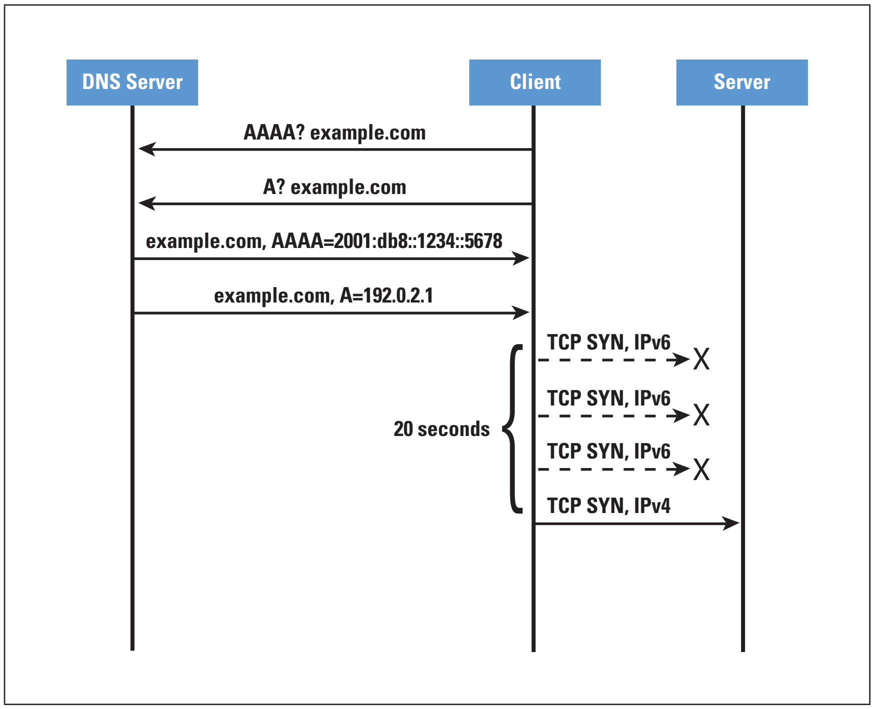

To work around that problem, Google implements a white list of DNS servers that it will provide *AAAA* records for[8]. However, in its current incarnation, DNS white listing does not scale well because the *Internet Service Provider* (ISP) has to prove good IPv6 connectivity to Google, and then Google white lists the ISP's DNS servers to receive the *AAAA* records. The scaling problem is that there are thousands of ISPs around the world, and white listing and de-white listing them becomes a tiresome manual task for both ISPs and Google. Furthermore, if every content provider did DNS white listing, ISPs would have to work with several content providers in order to give value to the IPv6 network they have deployed to their subscribers! Content providers have started working together to consolidate requirements for DNS white listing and operate some sort of DNS white-listing service to slightly automate this process[5].

Yet, DNS white listing still does not guarantee a working IPv6 network or a fast IPv6 network, because there is not a direct relationship between good IPv6 connectivity and the DNS server of a user's ISP. Even with the best of intentions and network design, there will still be instances where an IPv6 path or IPv4 path is working when the other path is broken. The result will be excessive delays for IPv4-only clients or dual-stack clients, depending on what sort of breakage occurs.

This situation contributes to the user perception that the Internet, or the particular website being accessed, is "down." The user will visit a different site instead, possibly never returning to the site that was "down."

Happy Eyeballs

A different approach solves these problems. In this approach, rather than an application slowly trying to make a connection on IPv6 and then on IPv4, the application makes its connection attempts more aggressively over both IPv6 and IPv4. Initially, the connection attempts are made *simultaneously* (rather than serialized), in order to provide a fast user experience.

The simultaneous connection attempts consume a little extra network bandwidth and twice the connection attempts on the server. To reduce that chatter, a cache is also maintained to store the success or failure of connecting using IPv6 or IPv4. We nickname this approach "Happy Eyeballs"[1], because the "eyeballs" (users) are happier—their computer provides them immediate content, even if the network is suffering slow performance on IPv6 or IPv4 (Figure 2).

Figure 2: Dual-Stack Web Browser Implementing Happy Eyeballs

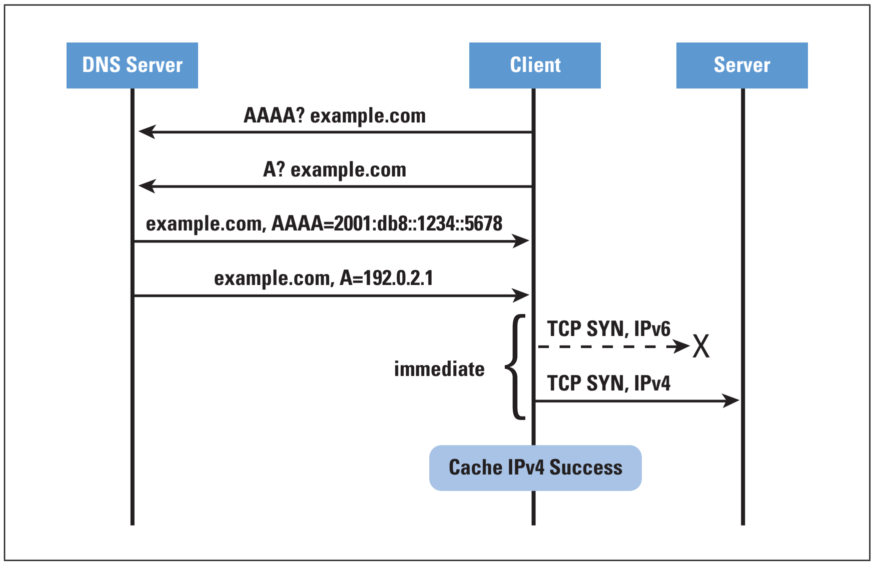

Obviously, sending a TCP SYN on both IPv6 and IPv4 doubles the number of connection attempts sent by the client. As discussed in [1], this chatter can be reduced by the application remembering if IPv6 (or IPv4) was successful in the previous connection attempt, and using that information for subsequent connection attempts. The sophistication of this cache is dependent on the memory (or disk) available, but even simple caching can be quite effective. When connecting to a new network (*third generation* [3G], different Wi-Fi network, or physical Ethernet), the connectivity of that new network can be determined and the cache of success or failure entirely or partially flushed, as necessary.

Thus, the doubling of connection attempts occurs only when connecting to a new network. Thereafter, initial connection attempts are delayed so that IPv6 (or IPv4) is tried first. But in all cases, significant user-noticeable delays are avoided when the IPv6 (or IPv4) is broken. The goal of Happy Eyeballs is to keep IPv6 enabled; that is, to make users unaware of IPv6 outages, so the user still visits IPv6-enabled websites without suffering any delay.

In this way, the user experiences a smooth migration from IPv4 to IPv6, and when necessary the fallback to IPv4 is almost immediate. This solution represents a significant improvement over today's web browsers. A drawback of this idea, however, is that it needs to be implemented in the application itself. Although it is a burden to upgrade those web browsers, there are only five major browsers[9], and the browsers receive the immediate benefit of the aggressive probing. Browsers are also commonly upgraded already for faster *JavaScript* engines and other new features.

Another idea to determine if IPv6 is working is to *ping* or send another simple request to an IPv6 resource on the Internet, and disable IPv6 on the host if that IPv6 request fails. This approach interferes with IPv6 traffic within the enterprise (which may be working fine, whereas IPv6 to the Internet is broken), and disabling IPv6 would break IPv6 features deployed in OSs (for example, *DirectAccess* in Windows or *Back to My Mac* in Mac OS X). An advantage of this approach is that if IPv6 is disabled, no application suffers the IPv6 outage and associated delay to fall back to IPv4.

New Transport: SCTP

Besides the problem of network layer protocol selection, a similar task can be performed at the transport layer. Maybe surprisingly, one more transport protocol exists besides TCP, namely *Stream Control Transmission Protocol* (SCTP). SCTP provides significant advantages over TCP, and it was designed with some of the lessons learned by TCP implementations and deployment[4] in mind.

Unlike IPv6 and IPv4, which have different DNS resource records (*AAAA* and *A*), we don't have a resource record to indicate that an application could, or should, use a different transport protocol. But even if we could indicate support for SCTP in DNS, the path might block it, reducing the usefulness of a DNS resource record. The path could be blocked by a NAT or firewall that expects only TCP or *User Datagram Protocol* (UDP).

Happy Eyeballs also describes a technique where a client can simultaneously try connecting using both TCP and SCTP. By necessity, this attempt is done entirely in the application, and the application would prefer the transport that responded faster and cache that information to reduce network chatter for subsequent connections to that server. This scenario is shown in Figure 3.

Figure 3: Client Implementing Happy Eyeballs for TCP/SCTP Selection

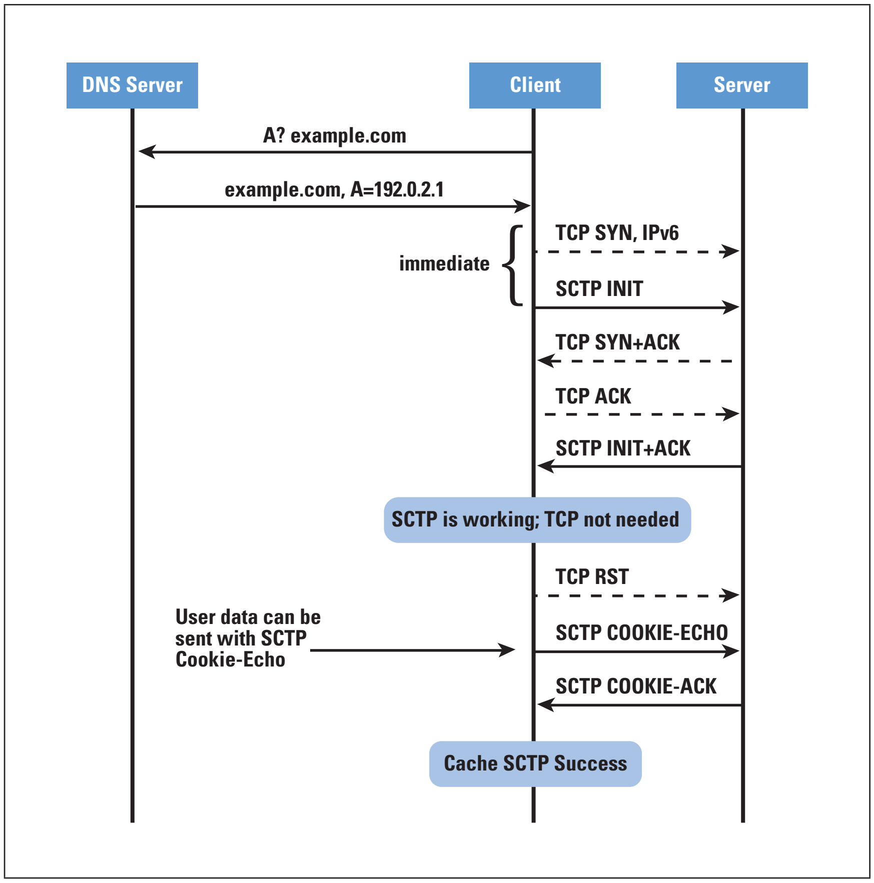

By combining the IPv6/IPv4 technique with the SCTP/TCP technique, a web browser running on a computer connected to a new dual-stack network sends four packets—an IPv4 TCP SYN, an IPv6 TCP SYN, an IPv4 SCTP INIT, and an IPv6 SCTP INIT. Based on the responses, it decides which transport protocol and which address family (IPv6 or IPv4) it prefers, and abandons the other connections. As described previously, connection information is cached for subsequent use to avoid consuming network bandwidth and server resources for subsequent network connections.

Conclusion

New technology aimed at improving user experience will be successful only if it meets expectations—an improved user experience. Because many companies are deriving all of their revenue from the Internet, any reduction in service means a loss of revenue. Thus, deploying new technology must not negatively affect the user experience. This article described one of the mechanisms that implementers can use to avoid negative effects on the user experience.

References

[1] Dan Wing, Andrew Yourtchenko, and Preethi Natarajan, "Happy Eyeballs: Trending Towards Success (IPv6 and SCTP)," Internet-Draft, Work-in-Progress, July 2009: `http://tools.ietf.org/html/draft-wing-http-new-tech`

[2] "Broken IPv6 clients," Lorenzo Colitti, June 2010: `https://sites.google.com/site/ipv6implementors/2010/agenda`

[3] "Google Trends": `http://www.google.com/trends?q=enable+ipv6%2C+disable+ipv6`

[4] P. Natarajan, "Leveraging Innovative Transport Layer Services for Improved Application Performance," February 2009: `http://www.cis.udel.edu/~amer/PEL/poc/pdf/NatarajanPhDdissertation.pdf`

[5] Carolyn Duffy Marsan, "Google, Microsoft, Netflix in talks to create shared list of IPv6 users," *Network World,* March 2010: `http://www.networkworld.com/news/2010/032610-dns-ipv6-whitelist.html`

[6] Tore Anderson, "IPv6 brokenness experiment, November results," November 2009: `http://lists.cluenet.de/pipermail/ipv6-ops/2009-December/002707.html`

[7] Igor Gashinsky, "IPv6 & recursive resolvers: How do we make the transition less painful?" March 2010: `http://www.ietf.org/proceedings/77/slides/dnsop-7.pdf`

[8] "Access Google services over IPv6": `http://www.google.com/intl/en/ipv6`

[9] "Usage share of web browsers": `http://en.wikipedia.org/wiki/Usage_share_of_web_browsers`

[10] R. Stewart, Ed., "Stream Control Transmission Protocol," RFC 4960, September 2007.

[11] Gunter Van de Velde, Ole Troan, and Tim Chown, "Non-Managed IPv6 Tunnels considered Harmful," July 2009: `http://tools.ietf.org/html/draft-vandevelde-v6ops-harmful-tunnels`

DAN WING has a B.S. in Computer Science from Central Washington University and has co-chaired the IETF's BEHAVE working group since 2006. He is a Distinguished Engineer at Cisco Systems, where he works on IPv6 transition technologies and has 30 patents issued or pending. E-mail: `dwing@cisco.com`

ANDREW YOURTCHENKO is a graduate of St. Petersburg Technical University in Russia, and has been in the networking industry since 1995. He is a Technical Leader at Cisco Systems in the network security area, and at IETF Andrew participates in the areas of security, TCP protocol, and IPv6 transition.
E-mail: `ayourtch@cisco.com`

Letter to the Editor

In response to "NAT++: Address Sharing in IPv4," in *The Internet Protocol Journal,* Volume 13, No. 2, June 2010:

Excellent article Geoff, so good I read it twice. While reading your article I was reminded of a recent experience that falls in the category of "unintended consequences." Since one of your situation descriptions was similar to the one I'm in, I thought I would relay my circumstance and experience and see if I can make my point.

A couple of months ago I signed up for an IPTV trial with my provider, and it was installed with a minimum of effort. The service is based on Cisco *Dial-on-Demand Routing* (DDR) and, of course, DSL service.

It worked fine for a couple of days; video feeds were good and all my computers and server worked just as before on a wireless network within my home. Then one day it appeared that I had lost *Domain Name System* (DNS) service, because I couldn't get name resolution to work but could route using the raw IPv4 addresses. So, I placed a trouble ticket and, of course, the provider's first request was to cold boot the DDR device and everything in the house, which I did. Sure enough, upon bringing all components back up (except one), everything was fine.

A day or two later I had to print something and powered on my HP 6510 wireless printer, printed what I needed to print, and then discovered I had lost DNS service again. I placed a trouble call and my provider came out and replaced the DDR device I went through the cold boot process (except one device) and everything was OK until I brought the printer online and the trouble returned. By now I had this nagging memory that wouldn't surface; something about that printer... With the printer powered off I rebooted the DDR, fired up SharkWire, and everything looked and worked OK.

Then I powered up the HP printer and saw another nagging memory; it immediately performed an *Address Resolution Protocol* (ARP) broadcast of the v4 address `169.254.65.206`—the famous black-hole address from RFC 3927[1]. Immediately after the ARP broadcast, the printer put out the normal *Dynamic Host Configuration Protocol* (DHCP) request and was assigned one from the *Network Address Translation* (NAT) pool.

That's when I stepped back from looking at the "trees" and gazed upon the "forest" and realized, with some embarrassment, that the public side (access side) was using a single IPv4 address with *Port Address Translation* (PAT) so the DDR box was blocking all the outbound PAT addresses attached to the single IPv4 address. I wrote down the details and e-mailed them to my provider, and had revised code pushed to the DDR the next day. Problem fixed.

All of this discussion leads me to ponder about other situations of "hard codes" in the network, either RFC-based or circumstance-based, that will falter with a switch to IPv6. Not in the core but in the customer networks. These unintended consequences could be many. Does HP run a dual stack for IPv4 and IPv6? I doubt it.

How can we get customers and vendors thinking about possible long-ago workarounds that they may have hard coded using IPv4? Any other RFCs out there like 3927? (It used to be easy when there were only a few hundred RFCs.) That could be the most expensive portion of the transition, verifying code ...

Keep up the good work; your articles make me think a lot and I really enjoy them. And, yes, I do use them for reference quite often.

Regards,

—Paul Dover
pdover@centeriem.com

[1] S. Cheshire, B. Aboba, and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses," RFC 3927, May 2005.

The author responds:

Thank you Paul for this anecdote and the important lesson behind it. Over some 30 years of intense development we've managed to accumulate a sizeable volume of technical specifications. Indeed, in October 2010 the RFC Editor published RFC 6068, and I'm not sure that any individual could claim a deep familiarity with every one of them, let alone claim to have a good understanding of their potential interaction. So when we look at various transitional technologies to sustain this industry through the next few years of attempting to support a comprehensive dual stack network in the face of the forthcoming hiatus of supply of IPv4 addresses, it should not come as a surprise when some devices or configurations fail in strange and unexpected ways, simply because they adhere to a technical standard that perhaps we've lost sight of in the flurry of generating new transitional technologies.

—Geoff Huston
gih@apnic.net

Fragments

Dr. Jianping Wu Receives Postel Award

The *Internet Society* (ISOC) recently awarded its prestigious *Jonathan B. Postel Service Award* for 2010 to leading Chinese technologist Dr. Jianping Wu for the pioneering role he has played in advancing Internet technology, deployment, and education in China and Asia Pacific over the last twenty years.

Dr. Wu's best-known contribution is the development of the *China Education and Research Network* (CERNET) which he designed and developed to be the first Internet backbone network in China. Created to establish a nation-wide advanced network infrastructure to support education and research among universities, CERNET has since become the world's largest national academic network. Since 1998, Dr. Wu has also devoted his time to the design and development of a large-scale native IPv6 backbone in China that now serves to connect over 200 universities and millions of users.

The Postel Award was established by the ISOC to honour individuals or organisations that, like Jon Postel, have made outstanding contributions in service to the data communications community. Commenting on its presentation to Dr. Wu, Lynn St. Amour, President and CEO of ISOC said: "Jianping Wu has dedicated his career in China to developing a broadly accessible Internet that brings people together. Twenty years ago, Dr. Wu recognized the importance and future impact of the Internet and the pivotal role it would play in terms of its impact on social reform, technology advancement and economic growth for China. He has worked tirelessly to bring his vision to life. As a result, the networks that resulted from his determination and hard work have played an important role in driving Internet development in China and have had a significant impact on the Internet worldwide."

ISOC presented the award, including a US$20,000 honorarium and a crystal engraved globe, during the 78th meeting of the *Internet Engineering Task Force* (IETF) in Maastricht, The Netherlands 25–30 July 2010.

DNSSEC Deployed in the Root Zone

On July 16, 2010 the U.S. Department of Commerce's *National Telecommunications and Information Administration* (NTIA) and the *National Institute of Standards and Technology* (NIST) announced the completion of an initiative with the *Internet Corporation for Assigned Names and Numbers* (ICANN) and VeriSign to enhance the security and stability of the Internet.

The announcement marks full deployment of a security technology—*Domain Name System Security Extensions* (DNSSEC)[1]—at the Internet's authoritative root zone, which will help protect Internet users against cache poisoning and other related cyber attacks.

"The Internet plays an increasingly vital role in daily life, from helping businesses expand to improving education and health care," said Assistant Secretary for Communications and Information and NTIA Administrator Lawrence E. Strickling. "The growth of the Internet is due in part to the trust of its users—trust, for example, that when they type a website address, they will be directed to their intended website. Today's action will help preserve that trust. It is an important milestone in the ongoing effort to increase Internet security and build a safer online environment for users."

"Improving the trustworthiness, robustness and scaling of the Internet's core infrastructure is an activity that lines up strongly with NIST's mission, and we have been contributing to design, standardization and deployment of DNSSEC technology for several years," said NIST Director Patrick Gallagher. "The deployment of DNSSEC at the root zone is the linchpin to facilitating its deployment throughout the world and enabling the current domain-name system to evolve into a significant new trust infrastructure for the Internet."

The *Domain Name System* (DNS) is a critical component of the Internet infrastructure. The DNS associates user-friendly domain names (for example, `www.commerce.gov`) with the numeric network addresses (for example, `170.110.225.168`) required to deliver information on the Internet, making the Internet easier for the public to navigate. The authenticity of the DNS data is essential to Internet use. For example, it is vital that users reach their intended destinations on the Internet and are not unknowingly redirected to bogus and malicious websites.

The DNS was not originally designed with strong security mechanisms, and technological advances have made it easier to exploit vulnerabilities in the DNS protocol that put the integrity of DNS data at risk. Many of these vulnerabilities are mitigated by the deployment of DNSSEC, which is a suite of *Internet Engineering Task Force* (IETF) specifications for securing information provided by the DNS.

A main goal of this action—DNSSEC deployment at the root zone—is to facilitate greater DNSSEC deployment throughout the rest of the global DNS hierarchy. While deployment of DNSSEC will protect Internet users from certain DNS-related cyber attacks, users must continue to exercise vigilance in protecting their information online.

ISOC Embraces DNSSEC

The *Internet Society* (ISOC) recently announced that it has deployed DNSSEC, a set of extensions to the DNS that provides a level of assurance, for its `isoc.org` domain. The announcement builds on an announcement by the *Public Interest Registry* (PIR) that they have implemented DNSSEC for the entire `.org` top-level domain.

"We are pleased to be among the first organisations in the `.org` top level domain to deploy DNSSEC, as DNSSEC provides an important building block for increasing user confidence in the Internet," said Lynn St.Amour, President and CEO of the Internet Society. "Implementing DNSSEC for the `.org` top-level domain is an important step in ensuring the global Internet serves as a trusted channel for communication and collaboration and we applaud the PIR's efforts in this area."

"DNSSEC acts like tamper-proof packaging to make sure that when you type in the website name of your bank you actually get the server IP address your bank wants you to use," said Leslie Daigle, Chief Internet Technology Officer of ISOC. "In this way, DNSSEC allows us to have more confidence in the online activities that are increasingly becoming a part of our lives at work, home, and school."

DNSSEC technology used today is the result of careful protocol engineering and standardization within the IETF; implementation by various DNS vendors; and operational trials by DNS operators. In addition to `.org`, DNSSEC is currently implemented by several country-specific top-level domains: Brazil (`.br`), Bulgaria (`.bg`), The Czech Republic (`.cz`), Puerto Rico (`.pr`), and Sweden (`.se`).

ISOC is a non-profit organisation founded in 1992 to provide leadership in Internet related standards, education, and policy. ISOC is the organisational home of the IETF. With offices in Washington, D.C., and Geneva, Switzerland, it is dedicated to ensuring the open development, evolution, and use of the Internet for the benefit of people throughout the world. For more information see: `http://isoc.org`

DNSSEC Fund Announced

In order to speed up the process of introduction a more secure global DNS infrastructure, the Netherlands-based charity *NLnet Foundation* has announced the creation of a global fund where open source projects can apply for grants to work on *Domain Name System Security Extensions* (DNSSEC) in their Internet applications.

DNSSEC is one of the key technologies for a safer Internet, as it allows the Internet user to know for sure that he or she is being sent to the right computer or service on the Internet. "If you type the name of your bank into a browser, you want to be sure that you are actually directed to a computer of that bank," said Michiel Leenaars, Director of Strategy at NLnet foundation. "Domain names are vital to the way we use the Internet, and without DNSSEC users are open to serious abuse."

DNSSEC provides a cryptographic seal of authenticity that gives real proof of the validity of the domain name you use when you visit a website, chat or send an e-mail. With DNSSEC you get a *chain of trust* from the root of the Internet to the service you want to connect to—opening the way for many new exciting opportunities for humans and computers to exchange information safely. DNSSEC is being gradually introduced worldwide.

The new fund will provide grants for reengineering important software to reliably work with DNSSEC. "The signing of the root through DNSSEC is a historical moment, but in a way it is only the beginning," said Leslie Daigle, Chief Internet Technology Office at the Internet Society. "Actual users will not fully benefit from protection in the more challenging situations as long as DNSSEC does not reach them." A great deal of work has already been done at the infrastructure level—most DNS servers such as *BIND, NSD* and *Unbound* now support the new technology. However, it will take a lot of work at the user level as well: operating systems, web browsers, e-mail servers, VoIP clients, and many other pieces of software need to be able to reliably work with DNSSEC.

"Every Internet user deserves to be protected by DNSSEC, yet currently almost no end user software is ready to take full advantage of the availability of DNSSEC," said Leenaars. "The IT community has a big responsibility in making sure that DNSSEC gets deployed across the board swiftly. We aim to accelerate the process significantly by putting some money on the table, and we invite other stakeholders to join us."

Since there are many applications and platforms that will require work, the NLnet Foundation is very open to cooperation with others as well as to targeted donations from interested stakeholders such as governments, registries and corporations.

The NLnet Foundation is a registered Netherlands charity with a long history of supporting Internet standardization. The foundation gained its capital from selling the first Dutch Internet Service Provider.

Potential applicants and collaborators can find more information at:
`http://nlnet.nl/dnssec`

See also:

[1] Miek Gieben, "DNSSEC: The Protocol, Deployment, and a Bit of Development," *The Internet Protocol Journal,* Volume 7, No. 2, June 2004.

[2] Torbjörn Eklöv, and Stephan Lagerholm, "Operational Challenges when Implementing DNSSEC," *The Internet Protocol Journal,* Volume 13, No. 2, June 2010.

[3] `http://www.dnssec.net/`

Call for Papers: Internet Privacy Workshop

The *Internet Architecture Board* (IAB), *World Wide Web Consortium* (W3C), *Internet Society* (ISOC) and the *Massachusetts Institute of Technology* (MIT) will hold a joint *Internet Privacy Workshop* on December 8 and 9, 2010 at MIT, Cambridge, Massachusetts on the question:

"How Can Technology Help to Improve Privacy on the Internet?"

Information about who we are, what we own, what we have experienced, how we behave, where we are located, and how we can be reached are among the most personal pieces of information about us. This information is increasingly being made more easily available electronically via the Internet, often without the consent of the subject. The question for the workshop therefore is: How can we ensure that architectures and technologies for the Internet, including the World Wide Web, are developed in ways that respects users' intentions about their privacy?

This workshop aims to explore the experience and approaches taken by developers of Internet including Web technology, when designing privacy into these protocols and architectures. Engineers know that many design considerations need to be taken into account when developing solutions. Balancing between the conflicting goals of openness, privacy, economics, and security is often difficult, as illustrated by Clark, et al. in "Tussle in Cyberspace: Defining Tomorrow's Internet," see:

`http://groups.csail.mit.edu/ana/Publications/PubPDFs/Tussle2002.pdf`

As a member of the technical community, we invite you to share your experiences by participating in this important workshop. Workshop participants will focus on the core privacy challenges, the approaches taken to deal with them, and the status of the work in the field. The objective is to draw a relationship with other application areas and other privacy work in an effort to discuss how specific approaches can be generalized.

Interested parties must submit a brief contribution describing their work or approach as it relates to the workshop theme. We welcome visionary ideas for how to tackle Internet privacy problems, as well as write-ups of existing concepts, deployed technologies, and lessons-learned from successful or failed attempts at deploying privacy technologies. Contributions are not required to be original in content.

Submitters of accepted position papers will be invited to the workshop. The workshop will be structured as a series of working sessions, punctuated by invited speakers, who will present relevant background information or controversial ideas that will motivate participants to reach a deeper understanding of the subject.

The organizing committee may ask submitters of particularly topical papers to present their ideas and experiences to the workshop. We will publish submitted position papers and slides together with a summary report of the workshop. There are no plans for any remote participation in this workshop.

To be invited to the workshop, please submit position papers to **`privacy@iab.org`** by November 5, 2010. More detailed information about the workshop, including further details about the position paper requirements, is available at:

`http://www.iab.org/about/workshops/privacy/`

We look forward to your input,

Bernard Aboba (IAB)
Daniel Appelquist (W3C)
Jon Peterson (IAB)
Karen Sollins (MIT)
Trent Adams (ISOC)
Karen O'Donoghue (ISOC)
Thomas Roessler (W3C)
Hannes Tschofenig (IAB)

Organizations Urged to Stop Delaying IPv6 Deployment

The *Number Resource Organization* (NRO), the official representative of the five *Regional Internet Registries* (RIRs) that oversee the allocation of all Internet number resources, recently unveiled the findings of a global, independent survey into organizations' IPv6 readiness. Funded by the European Commission and conducted by GNKS Consult and TNO, the study reveals that the majority of organizations are taking steps toward IPv6 deployment, as the IPv4 address pool continues to deplete rapidly.

IP addresses are critical for the operation of the Internet. Every Internet-enabled device needs an IP address to connect to the rest of the network. The biggest threat facing the Internet today is that less than 6% of the current form of IP addresses, IPv4, remains and the pool is likely to be completely depleted next year. This means that organizations need to adopt IPv6, the next-generation addressing protocol. There is a far larger pool of IPv6 addresses, allowing for more devices to connect to the Internet and helping to safeguard the sustainable growth of the Internet.

The survey, which polled over 1,500 organizations from 140 countries, highlights that organizations are increasingly aware of the need to deploy IPv6: approximately 84% already have IPv6 addresses or have considered requesting them from the RIRs. Only 16% of respondents have no plans to deploy IPv6 addresses.

The study also demonstrates that there are some misconceptions around the cost of adopting IPv6. Over half of all respondents noted that the cost of deployment was a major barrier for IPv6 adoption. While organizations might delay investing in IPv6, this may ultimately result in greater costs, with last-minute deployment and poor planning likely to increase the investment required.

Of the 84% of respondents that have requested IPv6 addresses or have considered doing so, three-quarters reported the need to stay ahead of competition as the main reason for IPv6 adoption. Half of these respondents also noted that a lack of available IPv4 space was a major driver for deployment. When asked about issues they had encountered when deploying IPv6:

- 60% cited the lack of vendor support as a major barrier for deployment. However, most of the latest hardware and software support IPv6. The RIRs are strongly urging organizations to check with their suppliers to ensure that the technologies they use are IPv6 compatible.
- 45% reported a struggle to find knowledgeable technical staff to support deployment. However, all five RIRs arrange technical training to facilitate an efficient IPv6 deployment, details of which can be accessed via the NRO website.

Fifty-eight percent of all organizations polled were ISPs. It is likely that respondents to this survey are further ahead in IPv6 deployment than ISPs overall, but all organizations should ensure that their ISP offers or plans to offer services over IPv6. Out of the polled ISPs:

- Approximately 60% already offer, or plan to offer within the next year, IPv6 to consumers.
- 70% already offer, or plan to offer within the next year, IPv6 to businesses.
- Only about 10% of polled ISPs have no plans to offer IPv6 to consumers or businesses.

Axel Pawlik, Chairman of the NRO, commented: "It's great to see that as we move toward complete IPv4 exhaustion, more organizations worldwide are waking up to the need to adopt IPv6 and are sourcing IPv6 addresses from the RIRs."

"Yet there is still a distinct lack of Internet traffic over the next addressing protocol, with not enough ISPs offering IPv6 services and 30% of ISPs saying the proportion of this traffic is less than 0.5%. It's critical that ISPs now take the next step in the global adoption effort by offering IPv6 services to their customers to help boost traffic over IPv6."

Per Blixt, Head of Unit in the Information Society and Medias at the European Commission, said:

"It's encouraging to see that so many organizations have made IPv6 adoption their priority. Still, as the Internet becomes increasingly important for global socio-economic development, it's critical that those who are still sitting on the fence act now on IPv6. Only by ensuring that all organizations adopt IPv6 can we ensure the sustainable growth of the digital economy worldwide."

This survey is a follow-up to a study conducted in 2009 amongst organizations in Europe, Middle East and parts of Central Asia, as well as Asia Pacific; however this year's survey polled organizations worldwide. The full research report is available at:

`http://www.nro.net/documents/GlobalIPv6SurveySummaryv2.pdf`

The NRO exists to protect the pool of unallocated Internet numbers (IP addresses and AS numbers) and serves as a coordinating mechanism for the five RIRs to act collectively on matters relating to the interests of RIRs. For further information, visit `http://www.nro.net`

The RIRs are independent, not-for-profit membership organizations that support the infrastructure of the Internet through technical coordination. There are five RIRs in the world today. Currently, the *Internet Assigned Numbers Association* (IANA) allocates blocks of IP addresses and ASNs, known collectively as *Internet Number Resources,* to the RIRs, who then distribute them to their members within their own specific service regions. RIR members include *Internet Service Providers* (ISPs), telecommunications organizations, large corporations, governments, academic institutions, and industry stakeholders, including end users

The RIR model of open, transparent participation has proven successful at responding to the rapidly changing Internet environment. Each RIR holds one to two open meetings per year, as well as facilitating online discussion by the community, to allow the open exchange of ideas from the technical community, the business sector, civil society, and government regulators. Each RIR performs a range of critical functions including: The reliable and stable allocation of Internet number resources (IPv4, IPv6 and *Autonymous System Number* resources); The responsible storage and maintenance of this registration data; The provision of an open, publicly accessible database where this data can be accessed. RIRs also provide a range of technical and coordination services for the Internet community. The five RIRs are:

AfriNIC:	`http://www.afrinic.net`
APNIC:	`http://www.apnic.net`
ARIN:	`http://www.arin.net`
LACNIC:	`http://www.lacnic.net`
RIPE NCC:	`http://www.ripe.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.

The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

*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*

Printed in the USA on recycled paper.