Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: September 3, 2012

P. McCann, Ed.
Huawei

March 2, 2012

Authentication and Mobility Management in a Flat Architecture
draft-mccann-dmm-flatarch-00

Abstract

Today's mobility management schemes make use of a hierarchy of tunnels from a relatively fixed anchor point, through one or more intermediate nodes, to reach the MN's current point of attachment. These schemes suffer from poor performance, scalability, and failure modes due to the centralization and statefulness of the anchor point(s). The dmm (Distributed Mobility Management) working group is currently chartered to investigate alternative solutions that will provide greater performance, scalability, and robustness through the distribution of mobility anchors. This document is an input to the dmm discussion. It outlines a problem statement for the existing mobility management techniques and goes on to propose (high-level) solutions to two of the most vexing problems: MN authentication and mobility management in a fully distributed, flat (non-hierarchical) access network. These two aspects are often treated separately in a layered architecture, but we argue there are important advantages to considering how these two functions can work in tandem to provide a simple and robust framework for the design of a wireless Internet Service Provider network.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 3, 2012.

Copyright Notice

Table of Contents

1. Introduction

Figure 1 depicts the hierarchical mobility management architecture that is being deployed by 3GPP LTE networks.

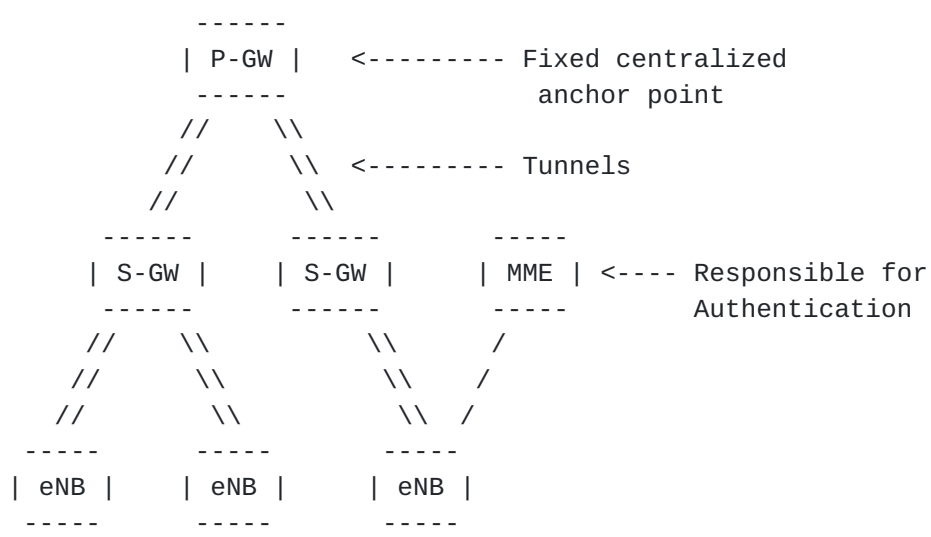

Figure 1: A typical hierarchical mobility management architecture.

This architecture is an evolution of the General Packet Radio System (GPRS) that was originally adopted for GSM systems. It is motivated in large part by two fundamental requirements:

- Keep the IP address (and therefore the anchor point) of the packet data session fixed for the life of the session; and,

- Re-use the existing legacy AKA authentication algorithm that was used for circuit voice.

These requirements were demanded by operators due to their desire to maintain control over services in the home network and to maintain their existing system of distributing user credentials in secure Subscriber Identity Modules (SIMs).

While these two requirements made sense for the operators that controlled standards decisions at the time, meeting them comes at somewhat of a cost. The use of a fixed P-GW without route optimization means that all packets have to traverse the chain of tunnels from anchor to MN, which could be very suboptimal if the MN is far away from the P-GW. The centralization of state in the P-GW (and to a lesser extent in the S-GWs) means that these nodes are scalability bottlenecks and that if one of them fails, all packet data sessions going through that node also fail. The re-use of a legacy symmetric secret key authentication protocol means that there

must be a round-trip to the home network to retrieve keying material upon initial attachment and every time the MN encounters a new MME. In addition to the performance impact, transport of secret keying material across inter-provider interfaces always carries some risk that the material will be compromised somewhere along the way. Note that keying material also needs to be transported from the MME to each eNB that the MN encounters.

This document examines the architectural implications of relaxing the above two requirements. In particular, we note that many MNs will not require a fixed IP address for the entire duration of their packet data session, as they will most likely be acting as clients and initiating short-lived connections to servers. It may be more important that such communication take the shortest path possible to reduce latency and load on the network. By making use of a routing protocol instead of a tunnel setup protocol for most mobility events, we can maximize the fault tolerance and compute the most optimal route for any packet from any vantage point in the network.

Second, we note that the hardware limitations that mandated the use of symmetric key algorithms for authentication are fading away. On a modern CPU, an elliptic curve public key cryptographic operation can be completed in well under 1 millisecond [1]. With the addition of low-cost cryptographic acceleration hardware [2], the battery impact of such an operation can be reduced even further. As CPU power is only increasing, we argue that it will be more important to reduce the number of messages and round-trips to the home network than to absolutely minimize the CPU consumption in the MN. Only a public key cryptosystem offers the ability to do this. With the creation of a new breed of authentication algorithms that can operate in one round-trip over the air, we can afford to perform a full re-authentication of the MN upon encountering each Access Router (AR), completely eliminating the need to transport secret keying material between infrastructure nodes.

The remainder of this document is structured as follows: Section 2 discusses the possibility of using a routing protocol for localized, network-based mobility management in a wireless access network. Section 3 introduces a possible public-key based authentication scheme that could be used for access authentication at each AR. Section 4 explores the synergy between authentication and mobility management, and explains how the new authentication algorithm could be embedded into Mobile IP for macro-mobility across domains. Section 5 is a list of work items for the IETF that will make this vision a reality. Finally, Section 6 and Section 7 give IANA and Security Considerations, respectively.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [3].

2. Mobility Management

Figure 2 gives a picture of a flat wireless Internet service provider network. Although ISP networks are usually structured in a hierarchy of layers such as Core, Aggregation, and Access routers, the connectivity between the routers is more mesh-like in nature and less rigidly hierarchical than the tunneling boxes shown in Figure 1.

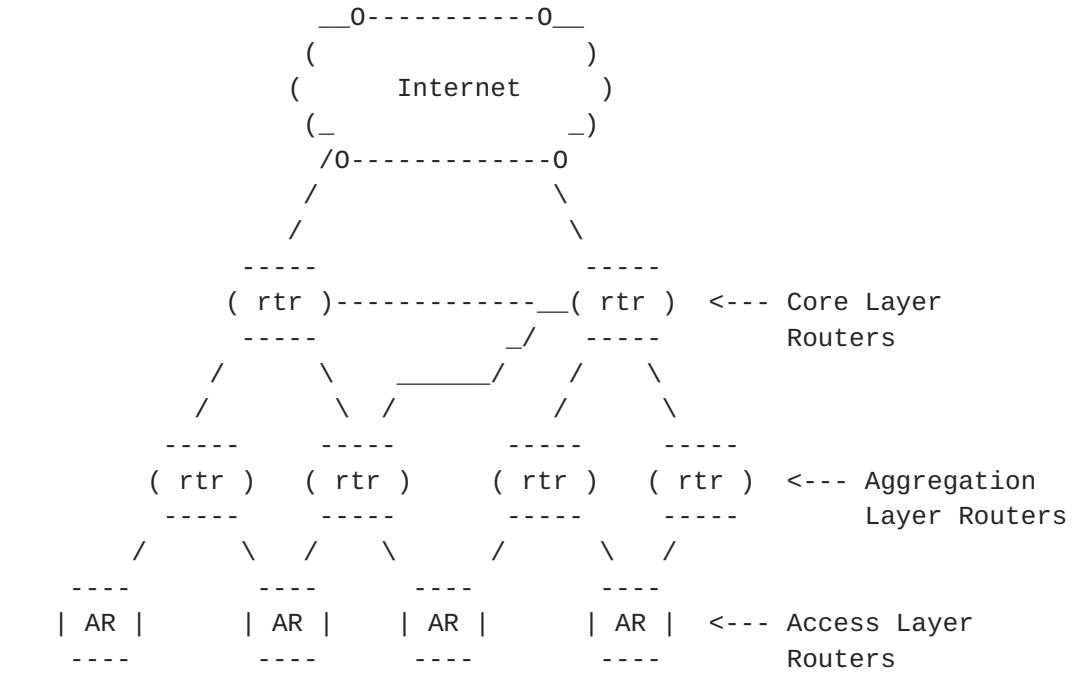

Figure 2: A flat wireless Internet service provider architecture.

The Access Routers (ARs) would be integrated with the radio link layer at the base stations. The ARs act as the first-hop routers for the MNs, and tunnels do not appear in the architecture until they are needed. Note also that each router can be connected to more than one router in the layer above, and can even be connected directly to some of its peer routers in the same layer. Except for the access layer, all of the routers in the network are standard off-the-shelf wireline routers running IBGP.

We assume that each AR has its own pool of addresses from which it can assign to mobile nodes and that these addresses are advertised

using IBGP to the upstream routers in the aggregation layer. We assume that all mobile nodes are authenticated upon attachment or re-attachment to a base station, and that the outcome of authentication is an exchange of hostnames (the base station learns the mobile node's hostname and vice-versa) bound to a master session key (MSK) shared between the mobile node and base station. Upon initial arrival in a given autonomous system, the mobile node is allocated an address (or a prefix) from the base station to which it is attached using ordinary mechanisms, e.g., DHCP. Then the mobile node updates its home DNS server to point from its hostname to the new address. The base station updates the reverse pointer in the in-addr.arpa or ip6.arpa space to point to the hostname it obtained when it authenticated the mobile node. Then, upon handoff, the target base station looks up the hostname received during authentication to determine whether the mobile node already has an address assigned from elsewhere in the autonomous system. If so, and if the hostname looked up in the reverse pointer is the same, it sends a BGP UPDATE message to all of its BGP peers containing the address (or the prefix) that was allocated to the mobile node. Packets are then routed appropriately to the new point of attachment in an optimal way. In the remainder of this section we describe the possible mechanisms in more detail.

2.1. Addressing Plan

The operator must define an addressing plan for the whole autonomous system. As a maximally-flat network, we assume that each base station will have its own designated pool of addresses from which it will assign to mobile nodes. To save space in the routing tables throughout the autonomous system, each pool should be a contiguous chunk of address space with a common prefix. Each base station acts as a BGP [4] router, originating UPDATE messages for the prefix(es) that it owns. The routers in the aggregation layer are configured as route reflectors [5] for the base stations they subtend (the base stations are route reflector clients). These routers are configured to aggregate the assigned address prefixes advertised by the base stations for the core routers above them, but will faithfully reflect all sub-prefixes advertised by any route reflector client to all other route reflector clients. Also, any sub-prefixes advertised by a client that are outside that client's pre-assigned range (known by configuration in the aggregation router) will also be reflected to the other clients and, if the prefix is outside the scope of the route reflector itself, propagated upward toward the core routers.

> On one popular BGP router platform, this would be accomplished with a combination of the "aggregate-address" command (without the "summary-only" option) and the "neighbor distribute-list out" command specifying that more-specific prefixes of the

known aggregate are to be suppressed to the non-client routers.

2.2. Handoff

Upon handoff within the same autonomous system, the mobile node is authenticated by the new base station. Given the mobile node's authenticated DNS name, the new base station takes several actions. First, it looks up the set of IP addresses associated with the hostname. It then makes a policy check on each IP address to see whether it is within the range of addresses managed by BGP in its local autonomous system. If so, it does a reverse lookup in the in-addr.arpa or ip6.arpa space (this space is controlled by the wireless ISP, unlike the forward mapping which is controlled by the mobile node) for each such IP address to ensure that some peer base station in its network did actually assign the IP address to the given name. If so, it originates a new BGP UPDATE message to its peers containing NLRI of the specific prefix (perhaps just a single address) that was assigned to the mobile node, with itself as the NEXT_HOP. It sets the LOCAL_PREF attribute to a 32-bit timestamp taken from its local clock (we assume that all base stations in an autonomous system have clocks synchronized to within 1 second). This will guarantee that the route is preferred over the same route that may have been advertised by a previously visited base station. The UPDATE will be sent to the parent routers in the aggregation layer, and will be reflected down to all other base stations in the same cluster. If the prefix was originally assigned by a peer base station in the same cluster, that is the extent of the update. Otherwise, the aggregation router propagates the update to the core layer which reflects it down to all other aggregation routers and from there it goes into all the base stations in the access layer.

Thus, when the mobile node moves within the same cluster, the messaging is confined to that cluster; however, when the mobile node crosses a cluster boundary, the update propagates through the larger cluster bounded by the route reflector above. If this is the core layer, then the update would be propagated throughout the autonomous system. This is necessary to ensure that optimal routes are created everywhere in the system. In general, there may be additional peer-to-peer links in the autonomous system; for example, directly between two neighboring base stations. Such a link would appear in the Interior Gateway Protocol (IGP, such as OSPF, EIGRP, or IS-IS) but would not be a BGP peering because the route reflectors take care of propagating BGP prefixes. Our scheme allows packets to make use of this route when appropriate; for example, a packet originated on one base station, destined for an IP address that is normally homed on the same base station but is being temporarily borrowed by a neighbor, would match the more-specific route to the neighbor listed as the NEXT_HOP in BGP and the recursive routing would forward the

packet over the direct link.

2.3. Address Management

The mobile node can therefore keep its address throughout the autonomous system if it wishes. When the address is nearing its lease expiration, the mobile node would send a unicast DHCPREQUEST to the DHCP server associated with the original base station to renew the lease. All base stations in the network must filter packets bound to IP addresses internal to the autonomous system to prevent abuse. In the case of DHCPREQUEST going to a remote base station, the current base station must add the DHCP Relay Agent Information Option [6] containing the mobile node's DNS hostname in the Agent Remote ID sub-option.

Keeping an address for a long period of mobility is sub-optimal due to the large amount of routing state that would be introduced. Our scheme is optimized for the case where the mobile node can periodically change its IP address to one that is more locally-appropriate. The BGP routing updates can provide a micro-mobility solution that hides the mobile node's movement from nodes outside the autonomous system and avoids frequent updates of its home DNS server. However, mobile nodes should keep track of which connections are using which addresses, and should periodically get new IP addresses from whatever base station to which they happen to be attached. Each IP address currently assigned to the mobile node should be registered to its home DNS server, with the most recently allocated listed first. Clients will therefore prefer the most recently allocated address for new connections.

Publishing the IP address assigned to a mobile node has security implications. Anyone who does a lookup can track the mobile node to the base station to which it was attached when it reserved the address. In general the use of an optimal route seems to be at odds with location privacy; if the mobile node needs location privacy, it should hide itself behind a proxy and only publish the proxy's IP address to the public DNS. Our scheme could function with pseudonyms assigned to mobile nodes by the visited network operator, but constructing such pseudonyms and allocating credentials to them is outside the scope of this document.

When a mobile node wants to release an address it should remove it from its home DNS server and send a DHCPRELEASE to the original assigning DHCP server. A DHCP server may have a policy that limits the number of times an IP address assignment may be renewed from a remote base station. This will force the mobile node to eventually release the address and optimize the routing tables.

The prefixes that we inject into the IBGP would most likely be full-length IPv4 addresses, although for IPv6 assignment of true prefixes would be more appropriate. All base stations in an autonomous system would need to agree on the prefix lengths they are assigning, and these prefixes would need to be on byte boundaries (for in-addr.arpa reverse lookups) or nybble boundaries (for ip6.arpa reverse lookups). The target base station would look up the mobile node's hostname and get back single IP addresses that are drawn from the prefixes and then do the reverse lookup on the containing prefix.

2.4. Macro-Mobility

The ability for any router in the access network to attract the packets destined for the MN can be used advantageously for macro-mobility as well as micro-mobility. Let's consider again the diagram from Figure 2, redrawn in Figure 3.

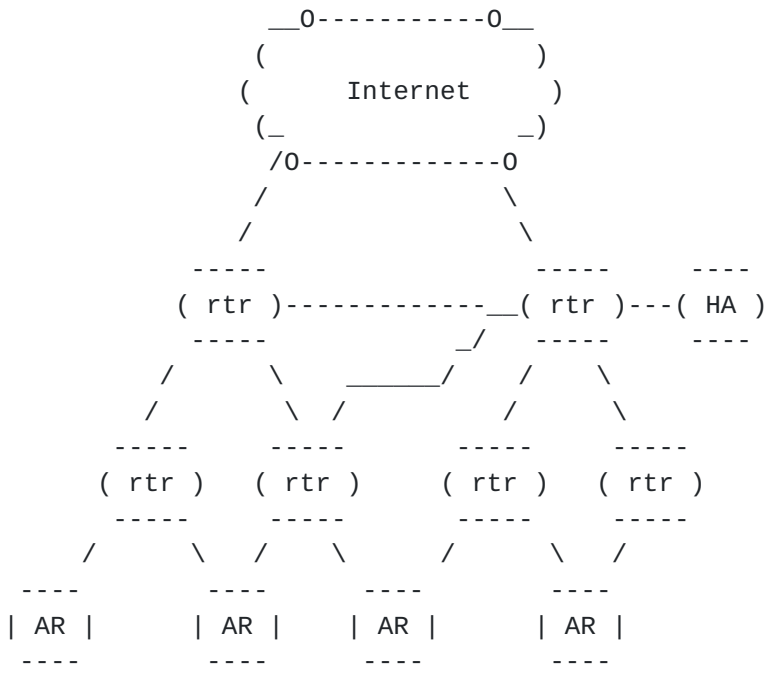

Figure 3: Home Agent support in a wireless Internet service provider architecture.

When the MN leaves the autonomous system completely, it may desire to keep sessions that were ongoing before it left. The Home Agent in the figure can be used for this purpose. Instead of using Gratuitous ARP or ND to attract the MN's packets, the HA can instead send a BGP UPDATE into the network to effect the routing of packets towards itself. This is a more powerful mechanism than ARP or ND because it can reach across multiple IP routing hops to install forwarding state that will route the packets in its direction.

The sending of a BGP UPDATE by the HA is triggered by an authenticated Registrationn Request or Binding Update. In this respect, the role of the HA can be compared to the role of any of the ARs that also must authenticate the MN before sending UPDATEs. We advocate a unification of the authentication protocol used for access and mobility signaling. The same set of credentials and secret keys can be used for both purposes, simplifying the network architecture and the node provisioning process. In the next section, we give a high level design for an authentication scheme that can be used.

3. Authentication

Recall in Section 2 we alluded to an authentication protocol that must run every time the MN encounters a new base station. To minimize the number of round-trips to the home network, we choose to base the authentication step on public key cryptography, namely an Elliptic Curve Diffie-Hellman exchange between the MN and the base station.

We note that the existing key exchange protocols such as IKE [7] and TLS [8] implement Perfect Forward Secrecy (PFS) by completing an ephemeral Diffie-Hellman exchange during the first round of messages between the communicating parties. Credentials are exchanged in a second round of messages. These multiple round-trips would introduce unacceptable overhead and latency in the mobile wireless environment. A key insight is that we don't necessarily need PFS if we are merely doing key exchange for purposes of authentication and integrity protection. A simpler protocol with one round-trip static Diffie-Hellman will suffice.

Perhaps the most difficult part of deploying a public key infrastructure is providing assurance that the public key obtained for the other party with which one wants to communicate does actually correspond to a private key known only to that other party. Key assurance can be provided through the use of certificates such as those defined by X.509 [9]. Usually, such certificates are exchanged in-band during the second round of a key exchange protocol. They must then be validated using e.g., the Online Certificate Status Protocol (OCSP) [10] or sometimes with an OCSP response attached ("stapled") to the same message that delivered the certificate. The exchange of certificates and OCSP information introduces additional overhead and possible round-trips to the authentication protocol.

In contrast, a new method of obtaining key assurance is currently being worked on in the DNS-based Assurance of Named Entities (DANE) working group [11]. While intended initially to support TLS, the protocol could be used for other purposes as well (e.g., S/MIME

[12]). Interestingly, some have proposed putting raw, bare public keys into the DNS records so that TLS can be run without the use of any certificates whatsoever [13]. It is this latter method of key assurance that we build on here.

Here we use the terminology of "peer" and "authenticator" as they are used in the Extensible Authentication Protocol (EAP) [14]. In our case the peer is the MN and the authenticator is the base station or Home Agent. We assume that the peer and authenticator are both named entities with DNS records containing the public portion of their keys. All such DNS records are protected with DNSSEC.

The peer and authenticator must discover each others' names and obtain the public keys corresponding to those names. There are several methods for how the peer might learn the authenticator's name and public key:

- The authenticator broadcasts its name and public key in system overhead messages.

- The authenticator unicasts its name and public key to the peer in an LTE Non-Access Stratum (NAS) message.

- The authenticator inserts its name and public key in the readable string portion of an EAP Identity Request and/or after the null terminating character.

- The peer somehow learns the DNS name of the authenticator and looks up the authenticator's key in the DNS using DNSSEC over an existing connection to the Internet prior to attaching to the authenticator.

In the first three methods, the peer may obtain assurance that the key belongs to the given name by making a DNS query as its very first action upon obtaining Internet access through the authenticator.

We assume that the authenticator has access to the Internet and can retrieve the key of the peer when it is given only the peer's DNS name during the authentication process. Distributing keys out-of-band helps to reduce the size of the authentication messages.

The actual authentication process consists of a single message sent from the peer to the authenticator. The message could be embedded in a NAS message or an EAP Identity Response message destined to the authenticator. Upon receiving and validating this message, the authenticator is able to derive a Master Session Key (MSK) which will be securely bound to the pair of DNS names given by both sides. The message from peer to authenticator would be protected with an HMAC

using the MSK derived by the peer. Upon validating the authentication message, and if requested by the peer, the authenticator may immediately begin the mobility management process outlined in Section 2. The authenticator may in parallel send a NAS message or an EAP Success message indicating successful authentication. The NAS message may be a Security Mode Command message that initializes the over-the-air integrity protection and encryption. The EAP Success message could trigger a lower layer key handshake as specified by IEEE 802.11i [15].

The derivation of the MSK is depicted below in Figure 4.

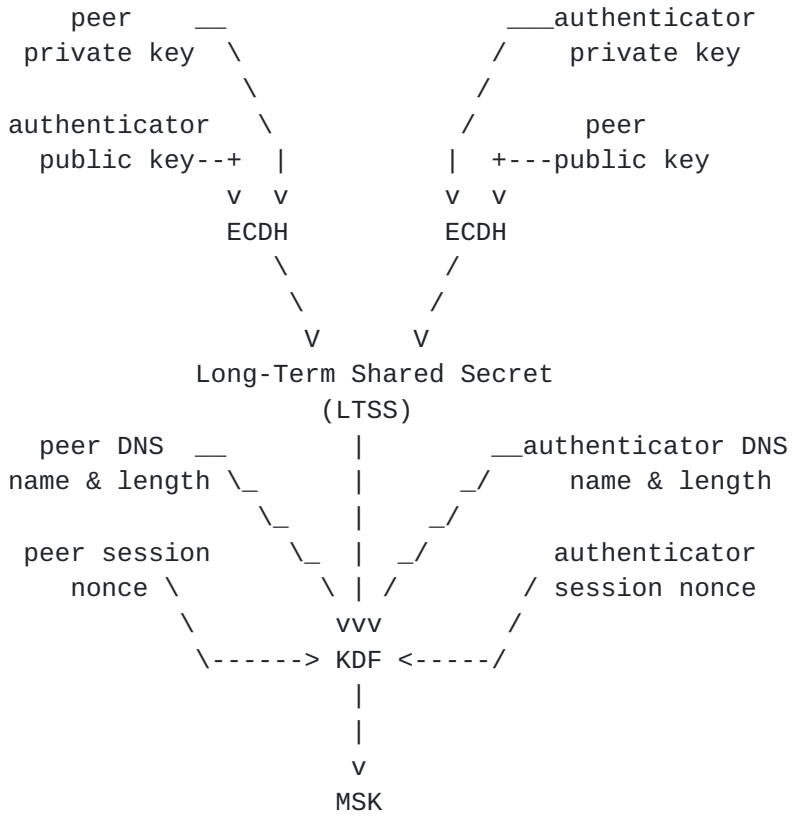

Figure 4: The key derivation hierarchy for authentication and key agreement in one round-trip.

In the figure, we depict the two sides independently performing Elliptic Curve Cofactor Diffie-Hellman (as specified in Section 5.7.1.2 of NIST SP 800-56A [16]). Each uses its own private key and the public key of the other entity. Both sides should arrive at the same value which they use as a Long-Term Shared Secret (LTSS). The LTSS may be cached so that the expensive ECDH operation does not have to be repeated when the same peer accesses the same authenticator in the future.

Next, we derive the MSK using a Key Derivation Function (KDF), taking as input the LTSS, the identities of the two parties (the DNS names and their lengths) as well as a session nonce from each party. The KDF should also include a counter value (set to 1) and a unique string indicating which function is calling the KDF. This will make the key derivation compliant to Section 5.8.1 of NIST SP 800-56A [16].

The authenticator may send a session nonce along with its public key in any of the four ways outlined earlier; note that in the case of publication in the DNS, the authenticator's session nonce would actually be re-used by incoming sessions for a period of time. Session MSKs would still be independent due to the entropy added by the peer in its own session nonce and by the different LTSSs derived for different peers.

> If it turns out that it is unacceptable to re-use the authenticator's nonce for more than one session, we will need to put an authenticator session nonce into the response to the peer's single authentication message. This response would trigger both sides to recompute MSK and to use it going forward. This response message should be authenticated with an HMAC using a key derived from either the first or second MSK to avoid denial-of-service attacks.
>
> Actually, it might be good to have such a "re-MSK" message available to either side during the life of the session to enable re-fresh of the MSK.

The derivation of the LTSS and the execution of the KDF to generate the MSK should be carried out in a secure environment, and both private keys and the LTSS should be stored in the secure environment so that they cannot be accessed except by the authentication method. The MSK may also be kept in the secure environment and an interface provided to derive further keys from it; alternatively, the MSK may be distributed to the outside environment for subsequent use. Historically, the secure environment has been implemented inside tamper-proof hardware that is resistant to duplication ("cloning"); such hardware usually runs at a much lower clock speed than the general purpose CPU that is used for other computing tasks. Because the ECDH operation will require the support of the main CPU, we envision that hardware virtualization support on the main CPU can be used to create a secure environment for the generation, storage, and use of private keys and the LTSS.

4. Mobility Management and Authentication Working Together

As described in Section 2, it is the completion of the authentication step that indicates to the AR that the MN is authentic and that its traffic should be redirected to the new point of attachment. Upon initial attachment, the MN doesn't have any assigned IP address and must obtain one using DHCP. At the same time, the DHCP server should assign the name of a Home Agent that can be used by the MN when it leaves the area inside which a BGP UPDATE accomplishes the traffic re-routing for the given address. The HA can be strategically placed at the boundary of this region, introducing the least amount of latency once the MN puts it on the forwarding path. The MN can perform a DNS lookup on the HA name to retrieve the HA's public key and perform the derivation of an LTSS long in advance of needing the HA's services. Messages could be provided so that the MN and HA can develop an MSK without the HA sending a BGP UPDATE; this would avoid the need to derive an MSK later when the Registration Request / Binding Update is actually sent.

We need some way of indicating to the MN whether or not its old address(es) have been successfully re-routed or whether it needs to perform a Mobile IP Registration Request / Binding Update to receive its traffic. One way is to wait for the AR to send a Router Advertisement (RA). The RA should contain all of those prefixes that were successfully re-routed by the AR sending a BGP UPDATE. If any prefix is missing from this list, the MN should initiate the Mobile IP Registration/Binding Update for those that are missing. However, this may be too much overhead so it may be desirable to build in some indication at the link layer (e.g., NAS signaling) when some prefixes were not able to be re-routed.

Existing LTE networks enable the MNs to remain in the idle state for many mobility events. This is accomplished through the use of Tracking Area Lists, and the MN does not need to update its location as long as it is within a Tracking Area that is on the list it was last sent. We can also support this concept; however, packets destined to the mobile node would always be routed to the AR on which it was last authenticated. That AR would need to page the MN throughout the Tracking Area List that it previously sent to the MN, and the MN would need to attach to the currently serving AR and carry out authentication to obtain these packets. The BGP UPDATE would reach the old AR which would then forward the packets as normal. The Tracking Area Lists should be chosen to make a proper tradeoff between the frequency of re-authentication and the size of the paging areas, keeping in mind that the MN will need to re-authenticate itself to receive packets at the current location.

Caching of the LTSSs will play an important role in improving the

performance of our scheme. Each MN could retain the LTSS for many if not all of the ARs it has previously visited, and the ARs could retain the LTSS for many of the recently seen MNs. This makes the derivation of the MSK a very simple matter of exchanging nonces and running the KDF.

5. Workplan for IETF

The IETF should undertake the following:

1. In the DANE working group, add authenticator nonces to the DNS record format for bare public keys.

2. Define a way to run the authentication protocol in this document over EAP.

3. In the DMM working group, define a way to run the authentication protocol in this document over Mobile IP. This may or may not involve running EAP over Mobile IP.

4. When defining the authentication protocol either over EAP or MIP, define a flag that allows the MN to control whether mobility management is immediately invoked or not (i.e., allow for derivation of the MSK by both sides without necessarily invoking mobility management).

5. Define a new DHCP option that carries a Home Agent DNS name.

6. Write an applicability statement and implementation guide for the use of BGP to create host routes for host mobility.

6. IANA Considerations

This memo includes no request to IANA as of yet.

7. Security Considerations

TBD

8. Acknowledgements

9. Informative References

[1] Bernstein, D., "Speed reports for elliptic-curve cryptography".

[2] Gaubatz, G., Kaps, J., Ozturk, E., and B. Sunar, "State of the Art in Ultra-Low Power Public Key Cryptography for Wireless Sensor Networks", 2nd IEEE International Workshop on Pervasive Computing and Communication Security (PerSec 2005) , 2005.

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

[4] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[5] Bates, T., Chandra, R., and E. Chen, "BGP Route Reflection - An Alternative to Full Mesh IBGP", RFC 2796, April 2000.

[6] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

[7] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

[8] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[9] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.

[10] Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP", RFC 2560, June 1999.

[11] Hoffman, P. and J. Schlyter, "Using Secure DNS to Associate Certificates with Domain Names For TLS", draft-ietf-dane-protocol-16 (work in progress), February 2012.

[12] Hoffman, P. and J. Schlyter, "Using Secure DNS to Associate Certificates with Domain Names For S/MIME", draft-hoffman-dane-smime-01 (work in progress), August 2011.

[13] Wouters, P., Gilmore, J., Weiler, S., Kivinen, T., and H. Tschofenig, "TLS out-of-band public key validation", draft-wouters-tls-oob-pubkey-02 (work in progress), November 2011.

[14] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.

[15] "Draft Standard for Information Technology - Telecommunications and information exchange between systems - Local and metropolitan area networks Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications", IEEE 802.11-REVma, 2006.

[16] Barker, E., Johnson, D., and M. Smid, "Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography (Revised)", NIST Special Publication 800-56A, March 2007, <http://csrc.nist.gov/publications/nistpubs/800-56A/SP800-56A_Revision1_Mar08-2007.pdf>.

Author's Address

Peter J. McCann (editor)
Huawei
400 Crossing Blvd, 2nd Floor
Bridgewater, NJ 08807
USA

Phone: +1 908 541 3563
Email: peter.mccann@huawei.com